climate*prediction*.net

Predicting Climate Change Through Volunteer Computing

Nick Faull

Project coordinator

University of Oxford Department of Atmospheric Physics

climate*prediction*.net

The Goals:

- To harness the power of idle home and business PCs to help forecast the climate of the 21st century.
- To improve public understanding of the nature of uncertainty in climate prediction.

The Method:

- Invite the public to download a full resolution, 3D climate model and run it locally on their PC.
- Use each PC to run a single member of a massive, perturbed physics ensemble.
- Provide visualization software and educational packages to maintain interest and facilitate school and undergraduate projects etc.

Volunteer Computing

- A specialized form of “distributed computing” which is really an “old idea” in computer science -- using remote computers to perform a same or similar tasks
- Was around before '99 but took off with SETI@home
- SETI@home peak cap with 500K users about 1 PF = 1000 TF
- for comparison Earth Sim in Kyoto = 35TF max
- climate*prediction*.net (CPDN) running at about 60 TF (60K concurrent users each 1GF machine average, i.e. PIV 2GHz conservatively rated)
- Offers high CPU power at low cost (need a few developers/sysadmins to run the “supercomputer”)

CPDN challenges…

- Model is about 1 million lines of legacy Fortran (40MB src)
- Proprietary, licenced by UK MetOffice
  - distribute executable/binary form only
- Resolution used: 2.75x3.75 degrees (73 lat x 96 long)
- Typically run on a supercomputer (i.e. Cray T3E) or 8-node Linux cluster (minimum)
- Ported to a single-processor, 32-bit Linux box
- Original: Windows only, now also Mac OS X, Linux
- Intel Fortran Win & Linux, IBM XLF for Mac, soon Intel Mac
- Many validation runs made on single-proc/32-bit to compare to supercomputer 64-bit
- Current coupled model takes ~6 months to run on a P4/2GHz PC 24/7!

climateprediction.net

BBC Climate Change Experiment

- http://www.bbc.co.uk/climatechange
- Participants download a 160-year atmosphere-ocean coupled model experiment (1920 to 2080)
- Promoted as part of the BBC “Climate Chaos” season programmes & documentaries
- The “Meltdown” documentary featured the CPDN project and launched the BBC/CPDN experiment in February of 2006
- Still has a strong user community and participant base
- Results show slated for December 2006

climateprediction.net Screensavers

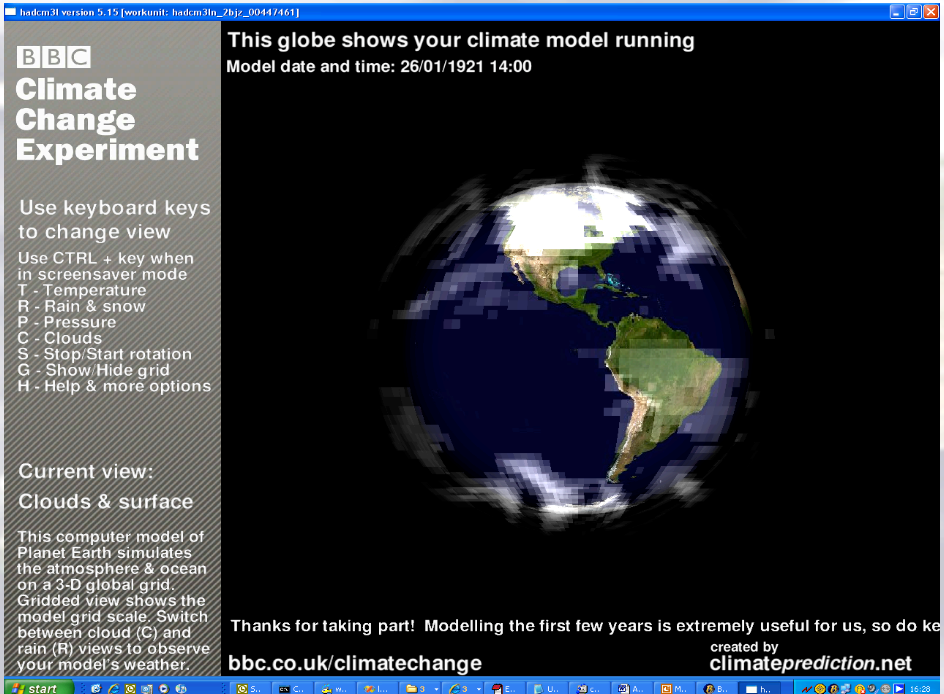

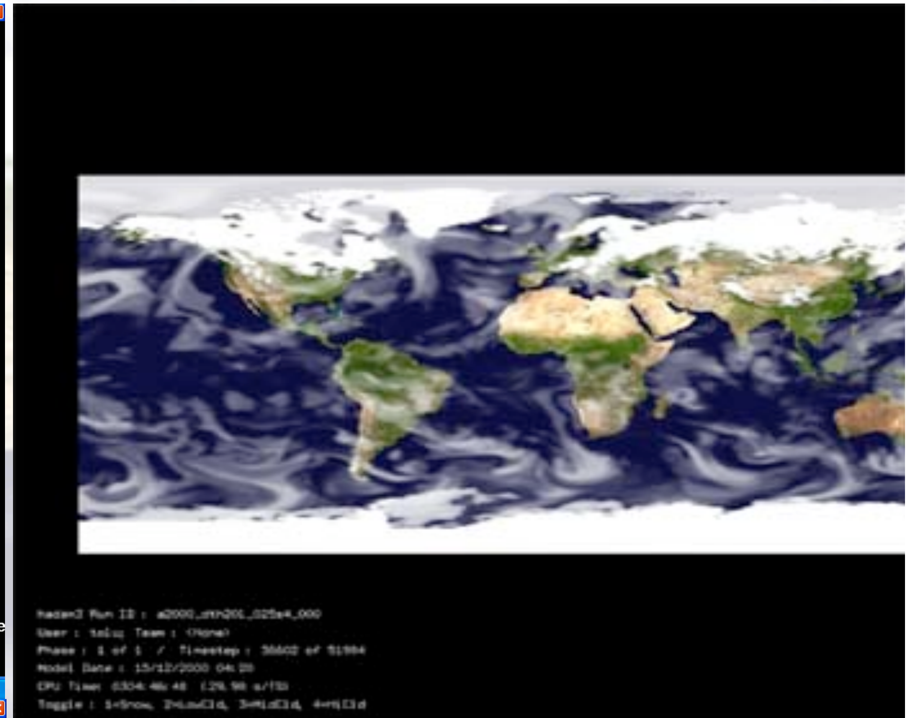

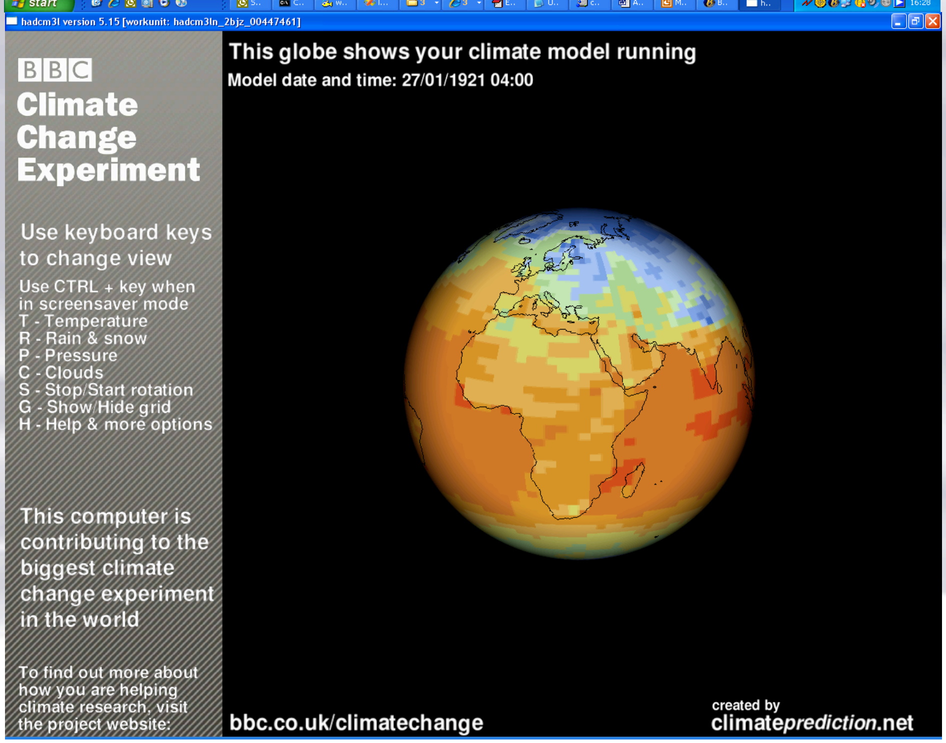

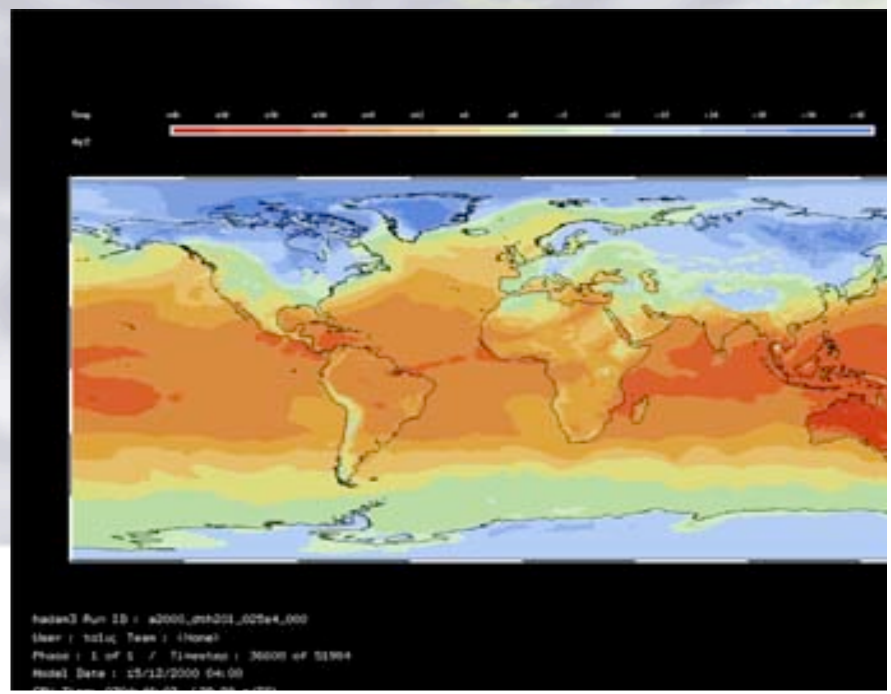

climate*prediction*.net for Educational Outreach

- CPDN has public education via the website, media, and schools as an important facet of the project
- Website has much information on climate change and related topics to the CPDN program.
- Open University (UK) offers a short course (S199) utilizing the climate*prediction*.net experiment (MS Windows client)
- Students hosted a debate on climate change issues, compared and contrasted their results, etc.

- Currently focused on UK schools, but as projects added and staff resources are gained plan to expand to other European schools and US schools

Students at Gosford Hill School, Oxon viewing their CPDN model

climateprediction.net Users Worldwide

>300,000 users total (90% MS Windows): ~60,000 active

~18 million model-years simulated (as of September '06)

~180,000 completed simulations

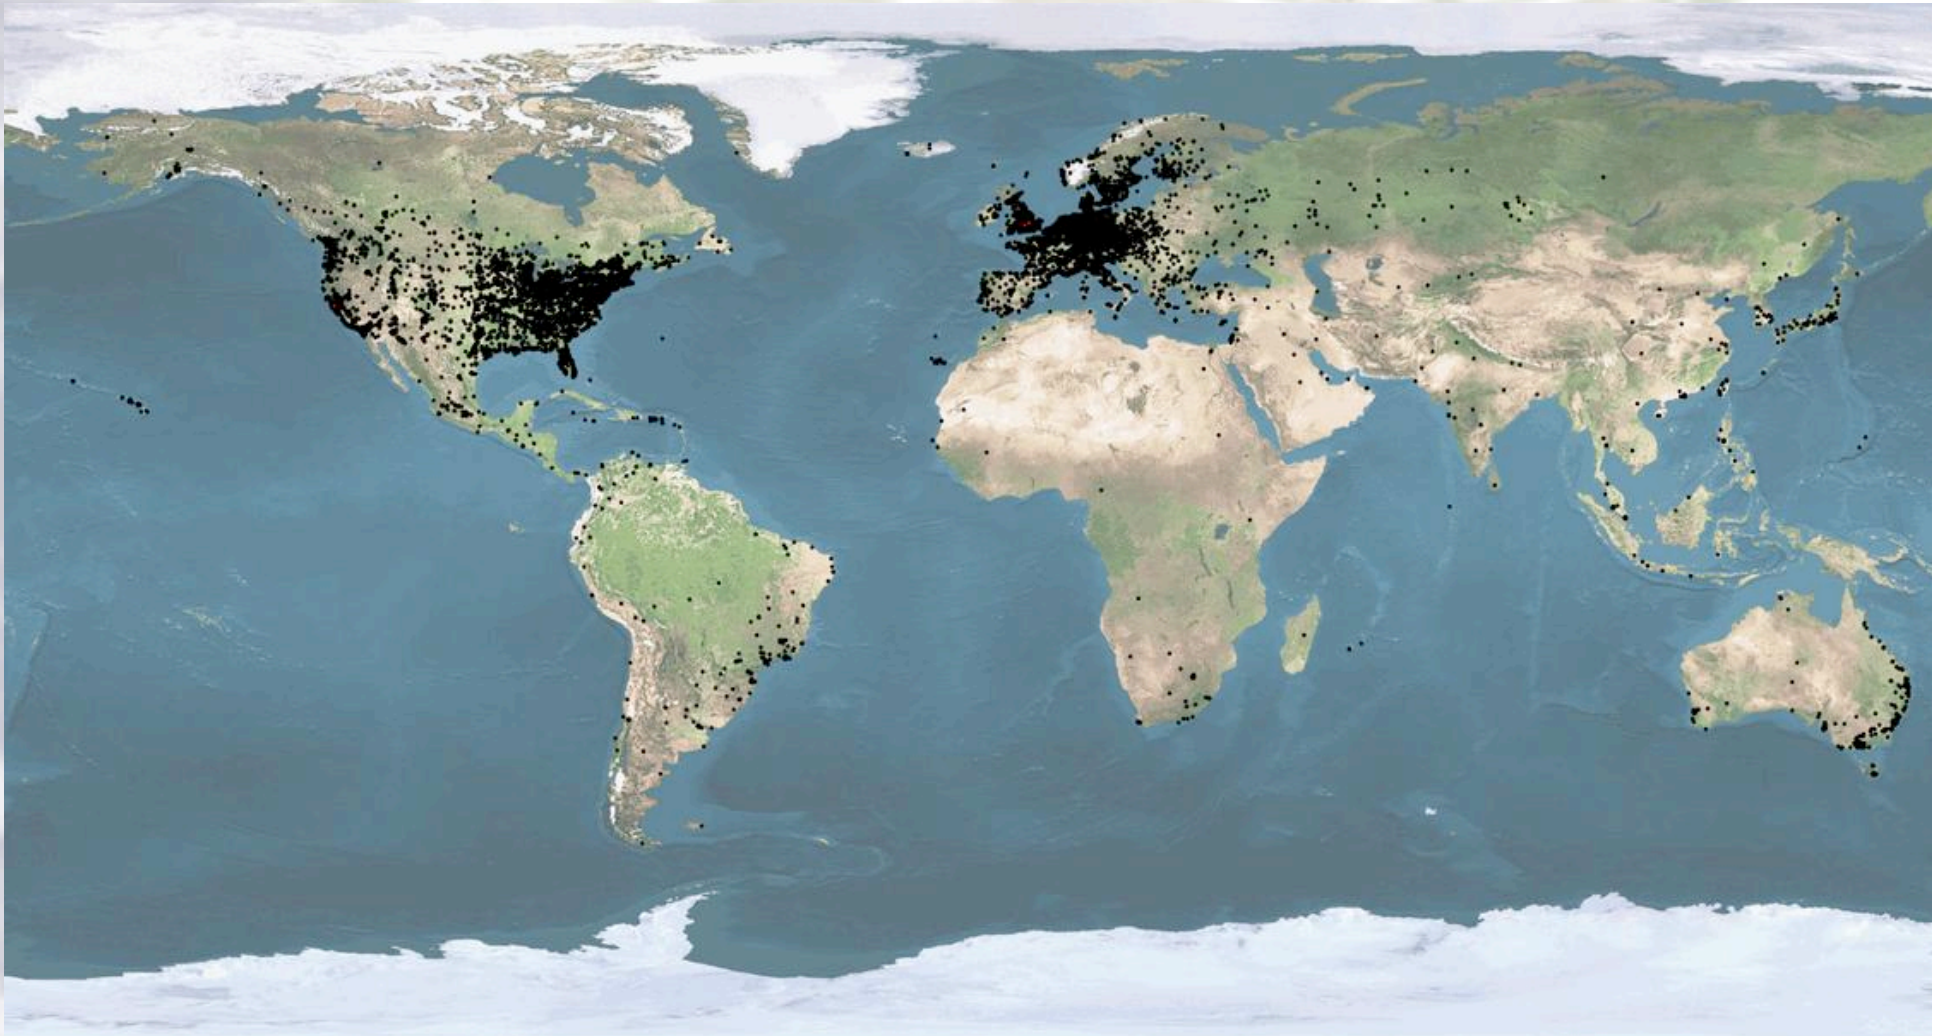

The world's largest climate modelling supercomputer!

(NB: a black dot is one or more computers running climateprediction.net)

Challenge – Keeping Users!

Ref: Carl Christensen, Tolu Aina, David Stainforth, *The Challenge of Volunteer Computing With Lengthy Climate Modelling Simulations*, **Proceedings of the 1st IEEE Conference on e-Science and Grid Computing**, Melbourne, Australia, 5-8 Dec 2005

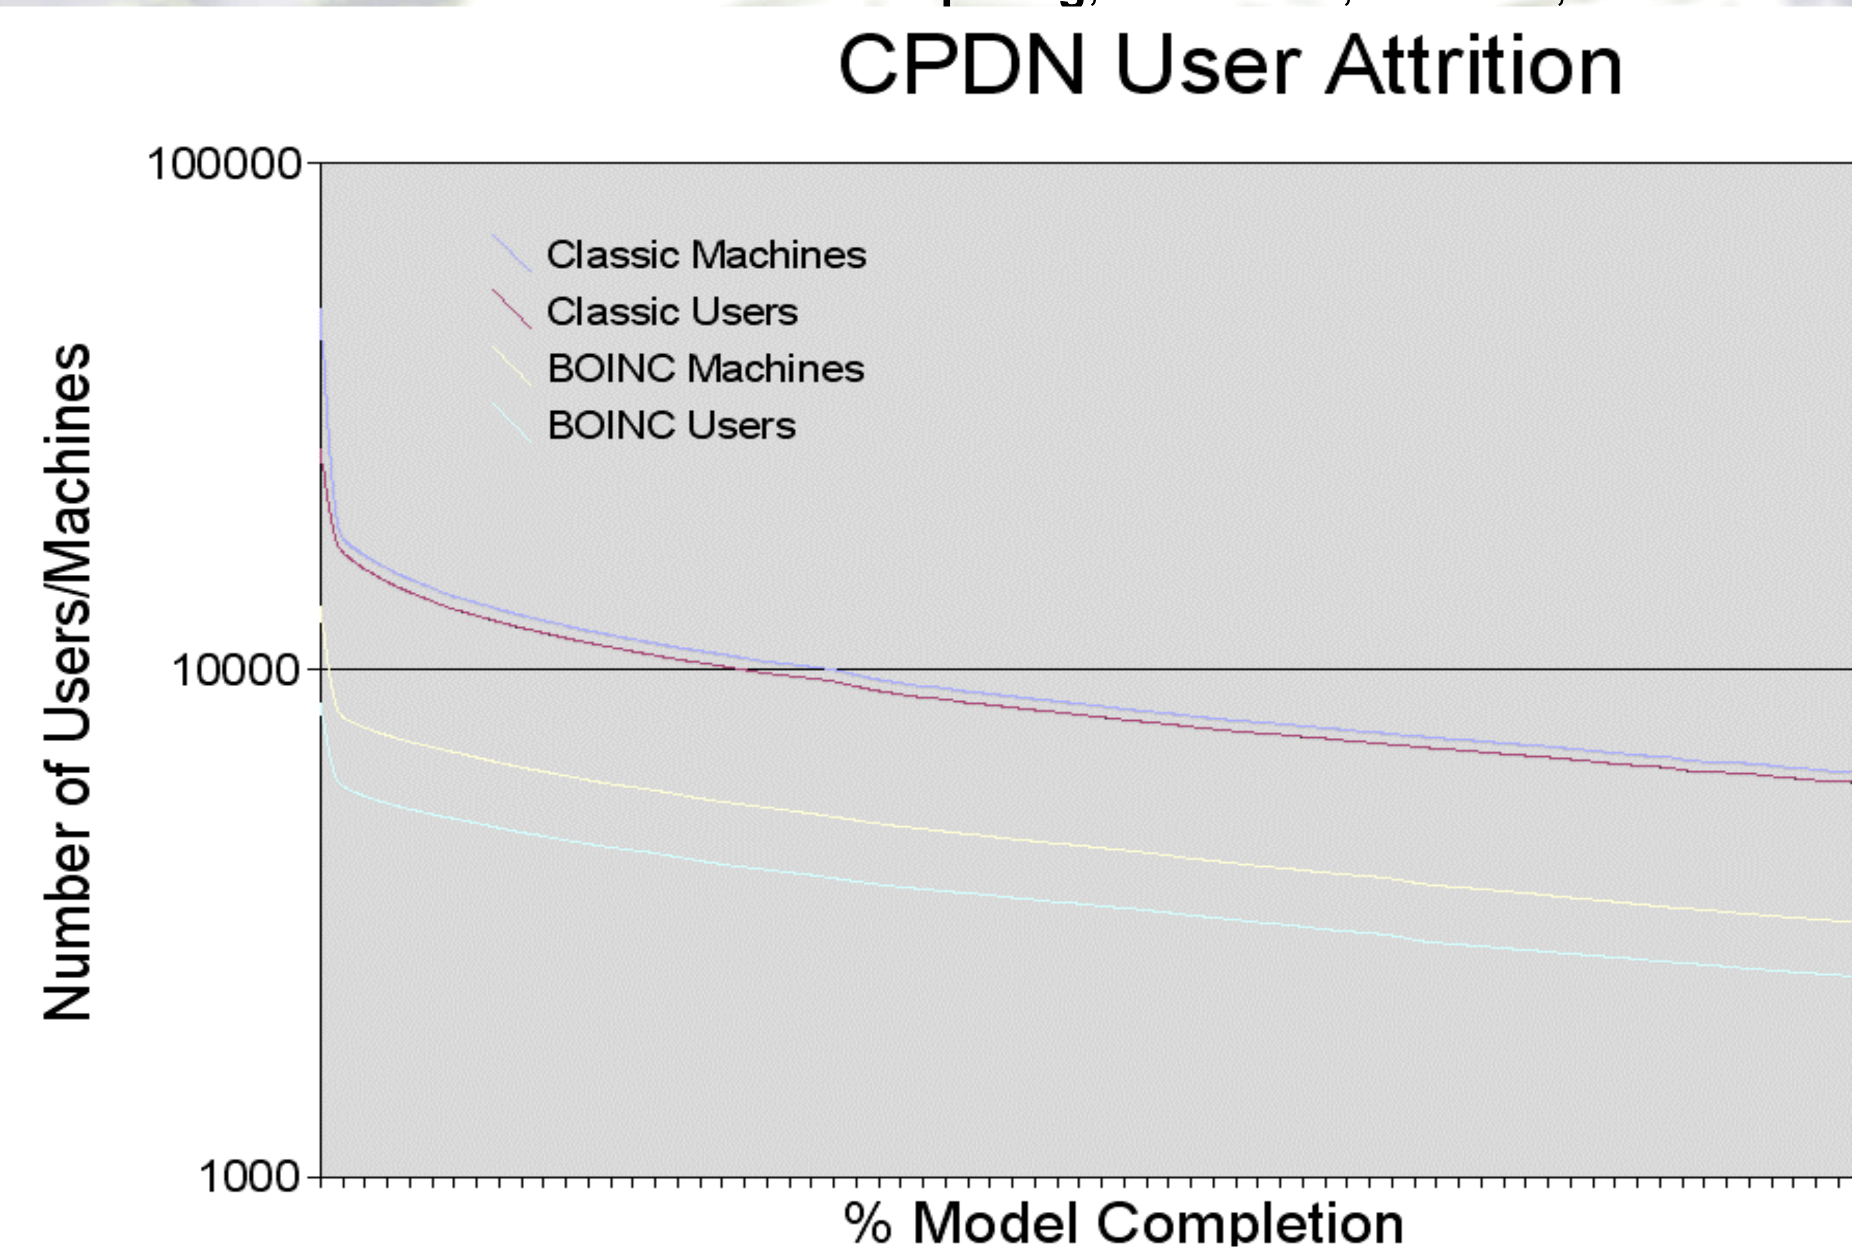

What goes into each model

- Natural Variability:
  The climate is chaotic with variations on timescales from minutes to centuries.
  Solution: Initial Condition Ensembles

- Forcing uncertainty:
  Changes due to factors external to the climate system e.g. greenhouse gas emissions (natural and anthropogenic), solar radiation etc.
  Solution: Scenarios for possible futures.

- Model uncertainty:
  Different models could be as could at simulating the past but give a different forecast for the future?
  Solution: Perturbed-Physics Ensembles

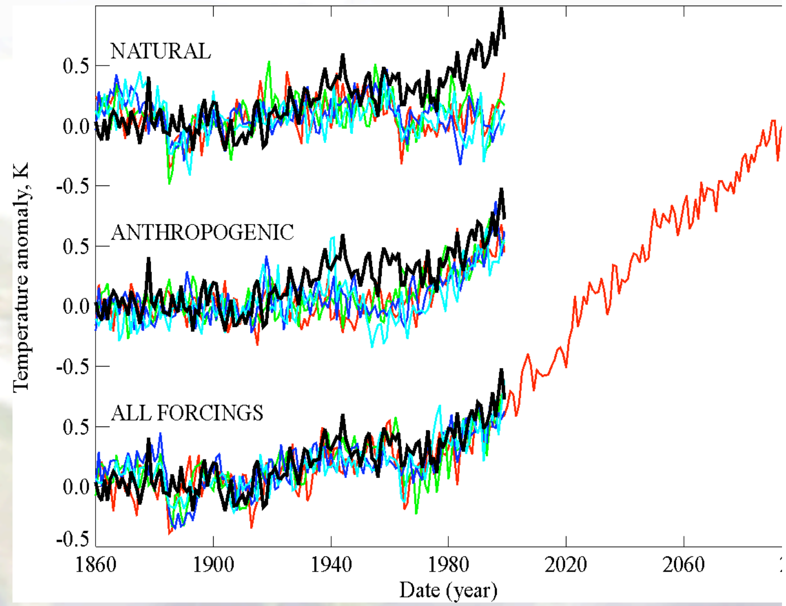

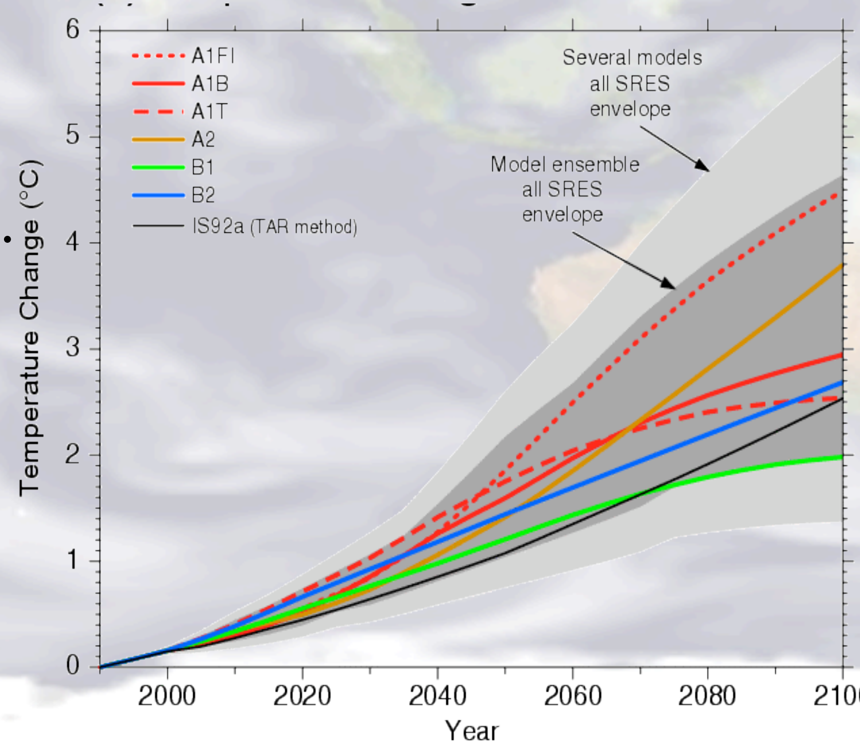

Example BBC Expt Run

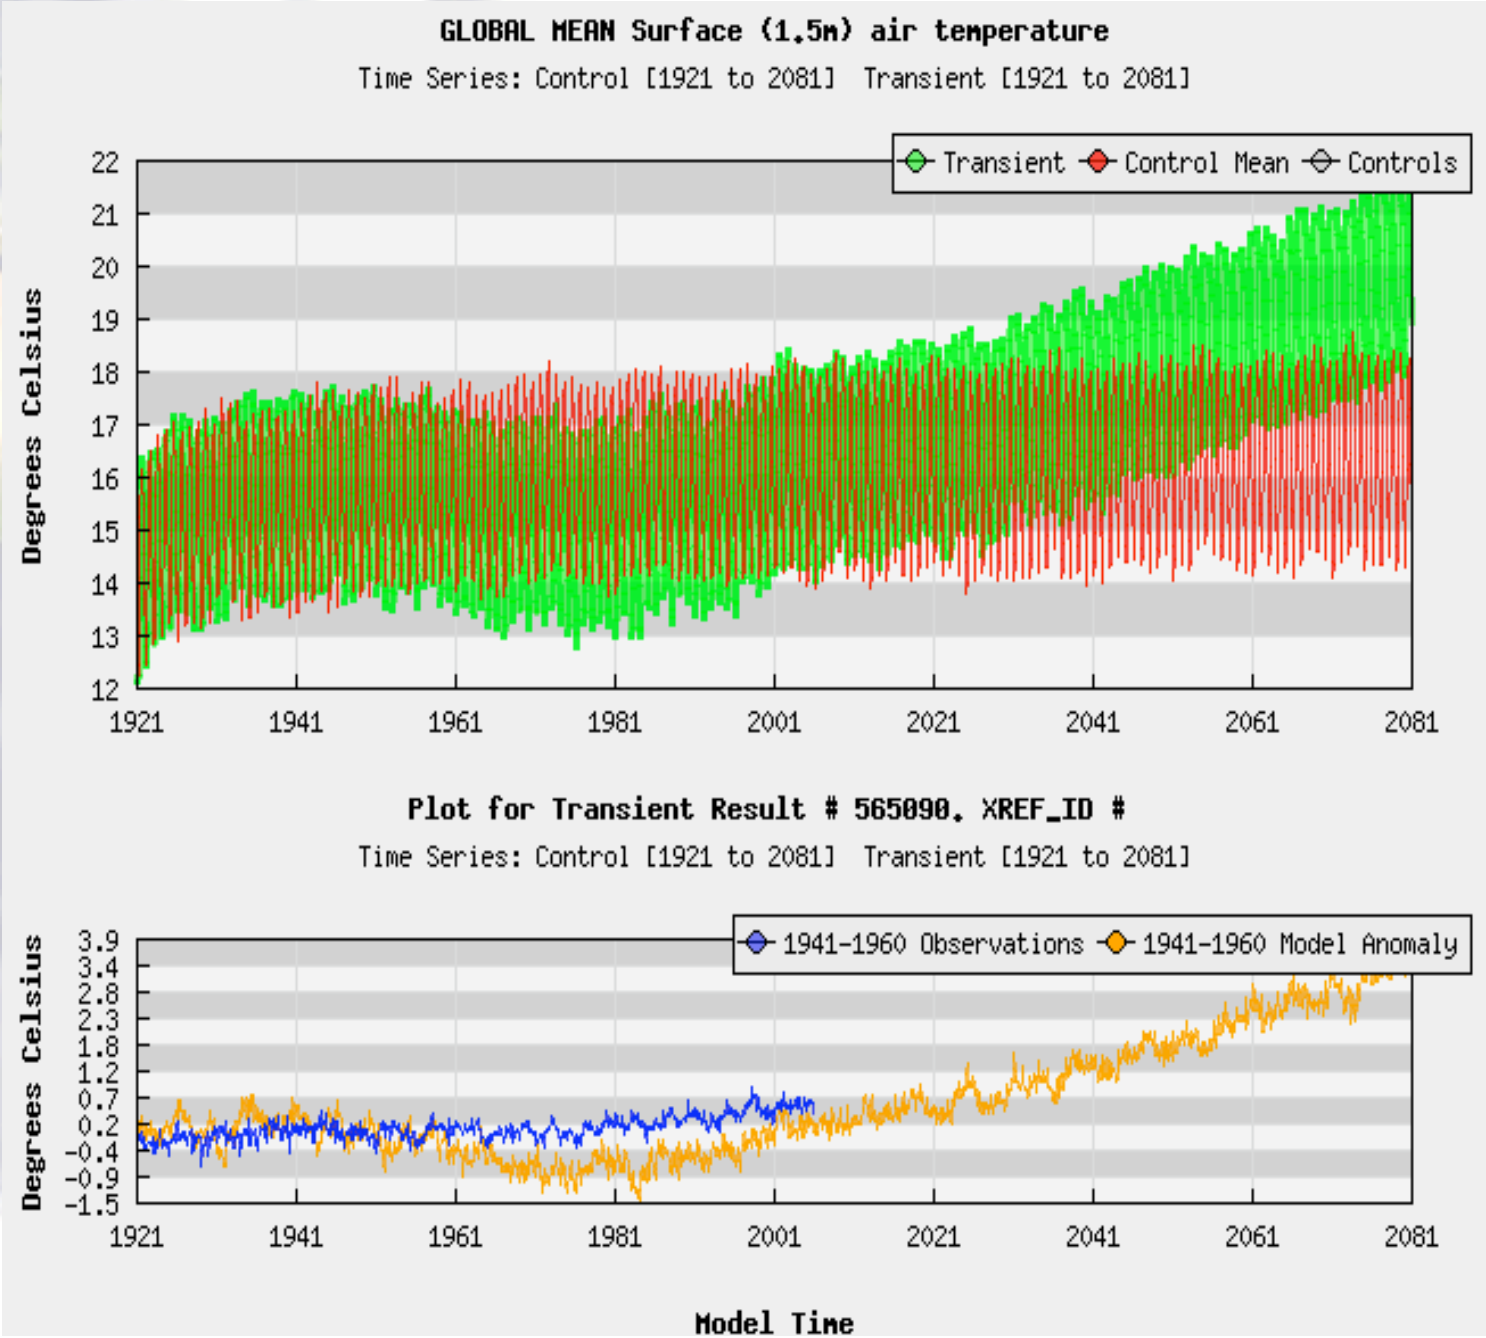

Plume Graph of 206 Runs
(Still one of the largest CM ensembles)

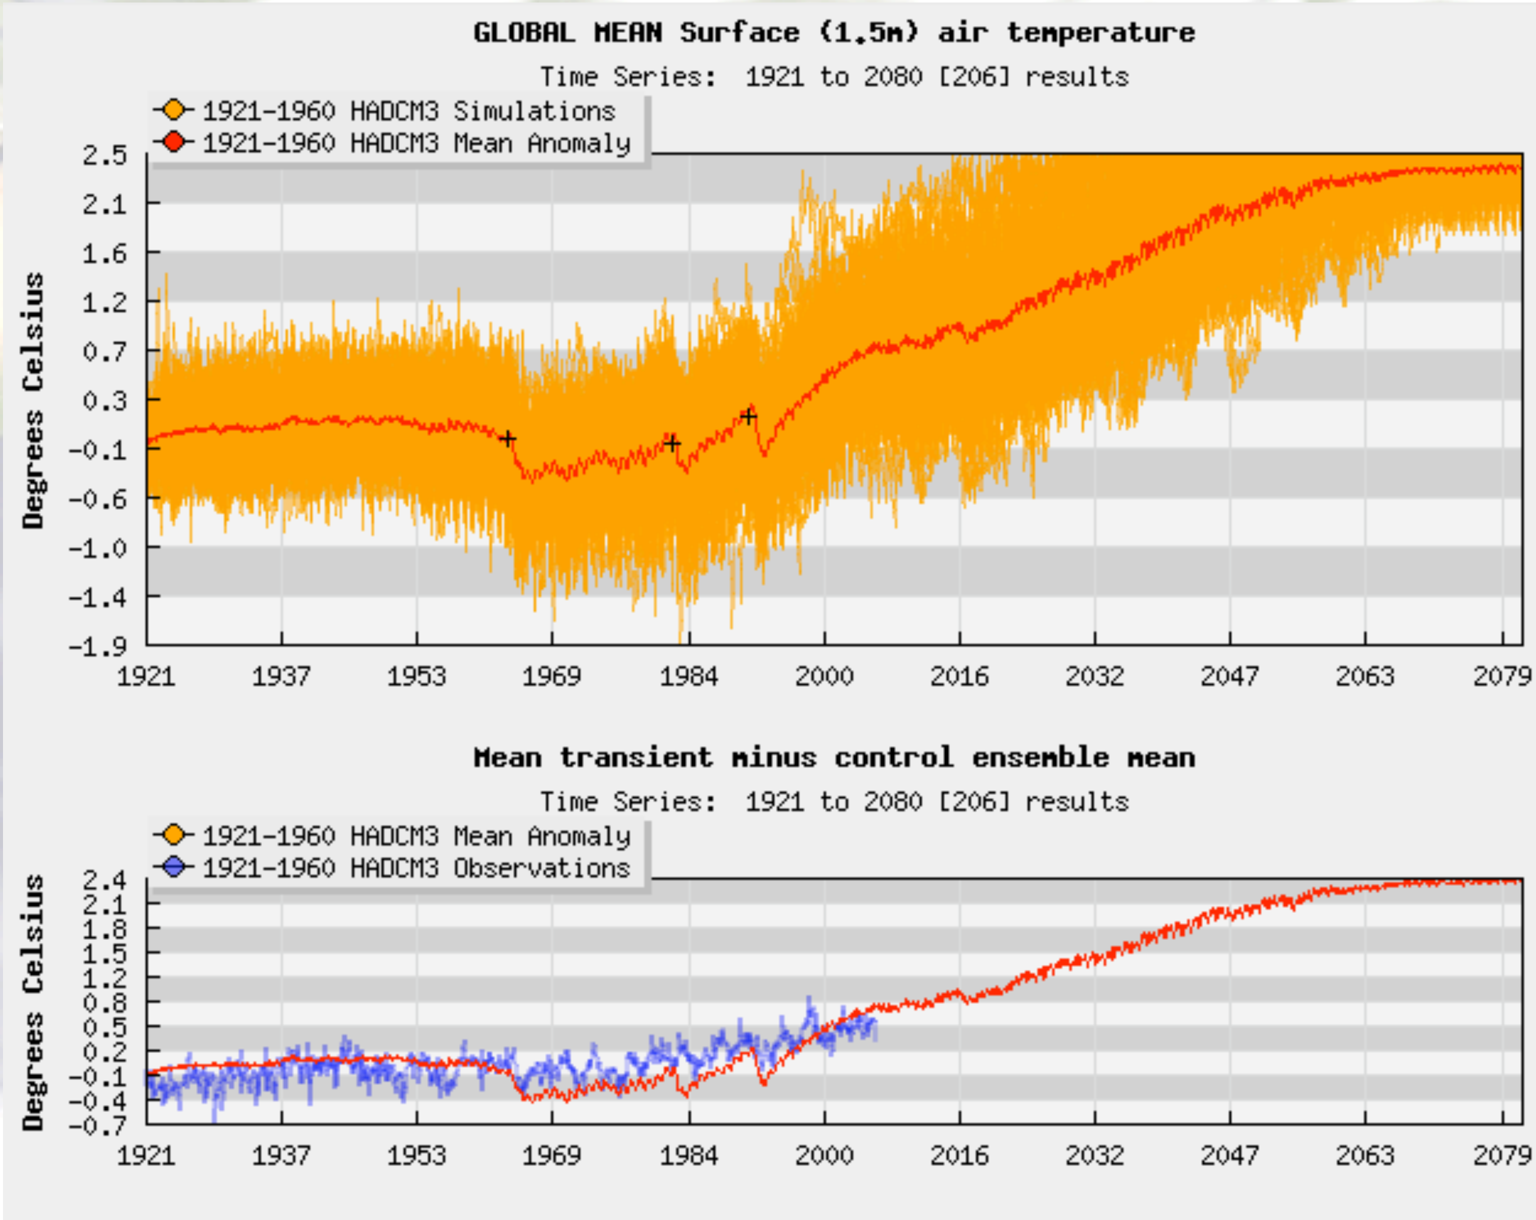

Results from first (simpler) experiment

- Launch in September 2003.
- Results published in January 2005 (Stainforth et al).

Results from our initial climate*prediction*.net experiment (Stainforth et al, 2005)

- Using simplified model ocean to keep runs short
  - 15-year calibration phase to compute ocean heat transport
  - 15-year control phase with pre-industrial CO_2 (280ppm)
  - 15-year $2xCO_2$ phase with CO_2 at 560ppm.
- Repeat with different initial conditions to average out noise and quantify sampling uncertainty

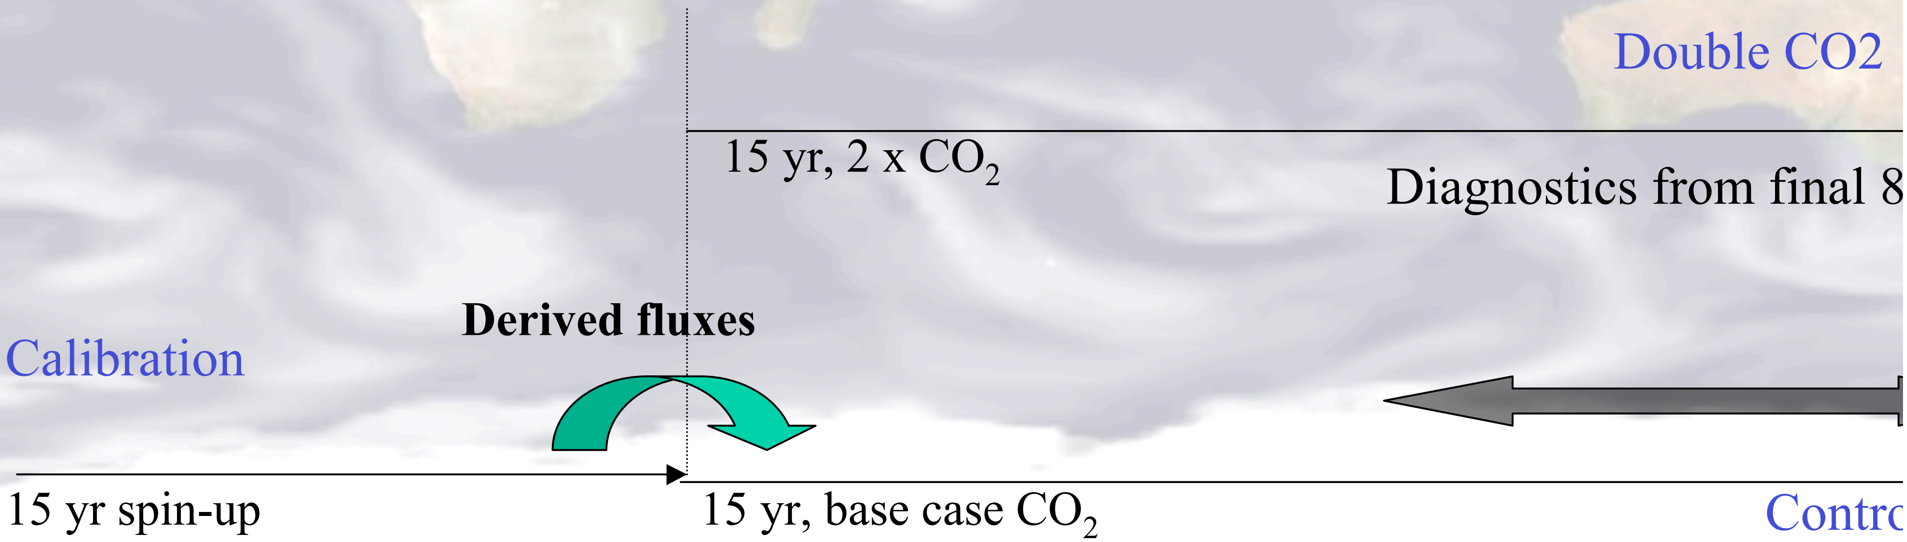

Extracting Climate Sensitivity

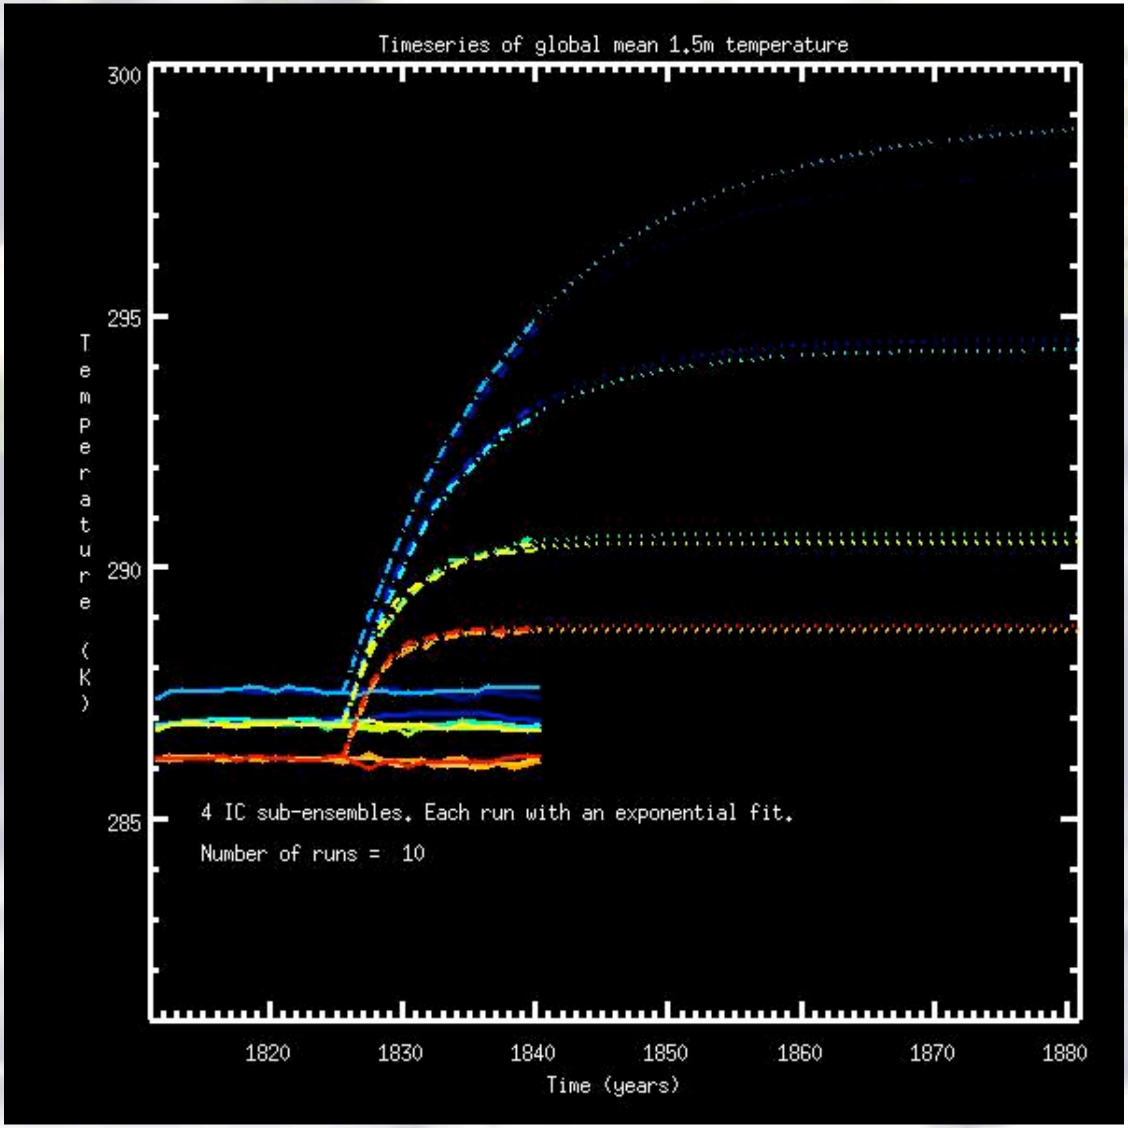

Climate sensitivity is th equilibrium global mean surface temperature change for a doubling of CO levels.

Parameter perturbations (initial results)

- Critical Relative Humidity [RHcrit]
- Accretion constant [CT]
- Condensation nuclei concentration [CW]
- Ice fall velocity [VF1]
- Entrainment coefficient (EntCoef).
- Empirically adjusted cloud fraction (EACF).

Parameter	Low value	Standard Value	High Value
VF1	0.5	1.0	2.0
CT	5.00E-05	1.00E-04	4.00E-04
RHcrit	0.6	0.7	0.9
CW (sea, land)	1.00E-04 2.00E-05	2.00E-04 5.00E-05	2.00E-03 5.00E-04
EntCoef	0.6	3	9
EACF	-	0.5	0.6

Frequency Distribution of Simulation

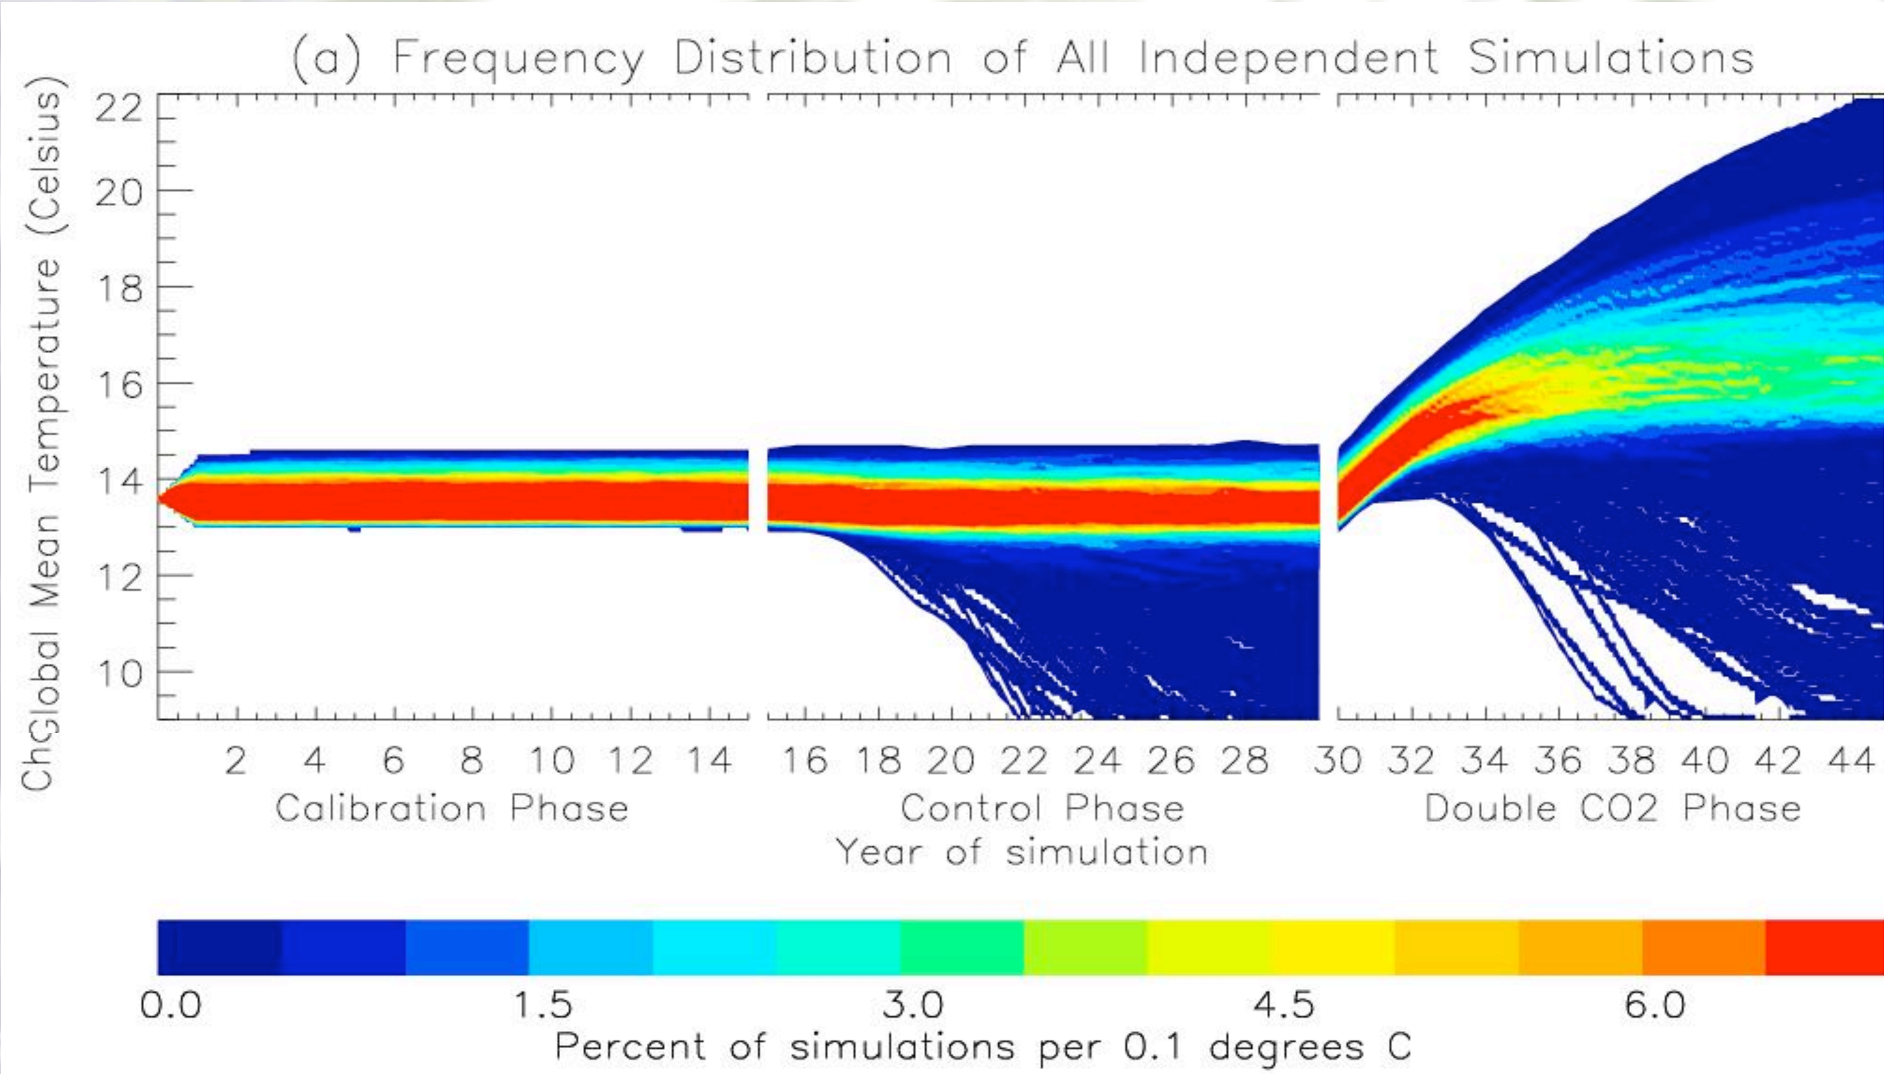

From Stainforth et al, Nature, 27 Jan ‘05

Frequency distribution, eliminating drifting control simulations

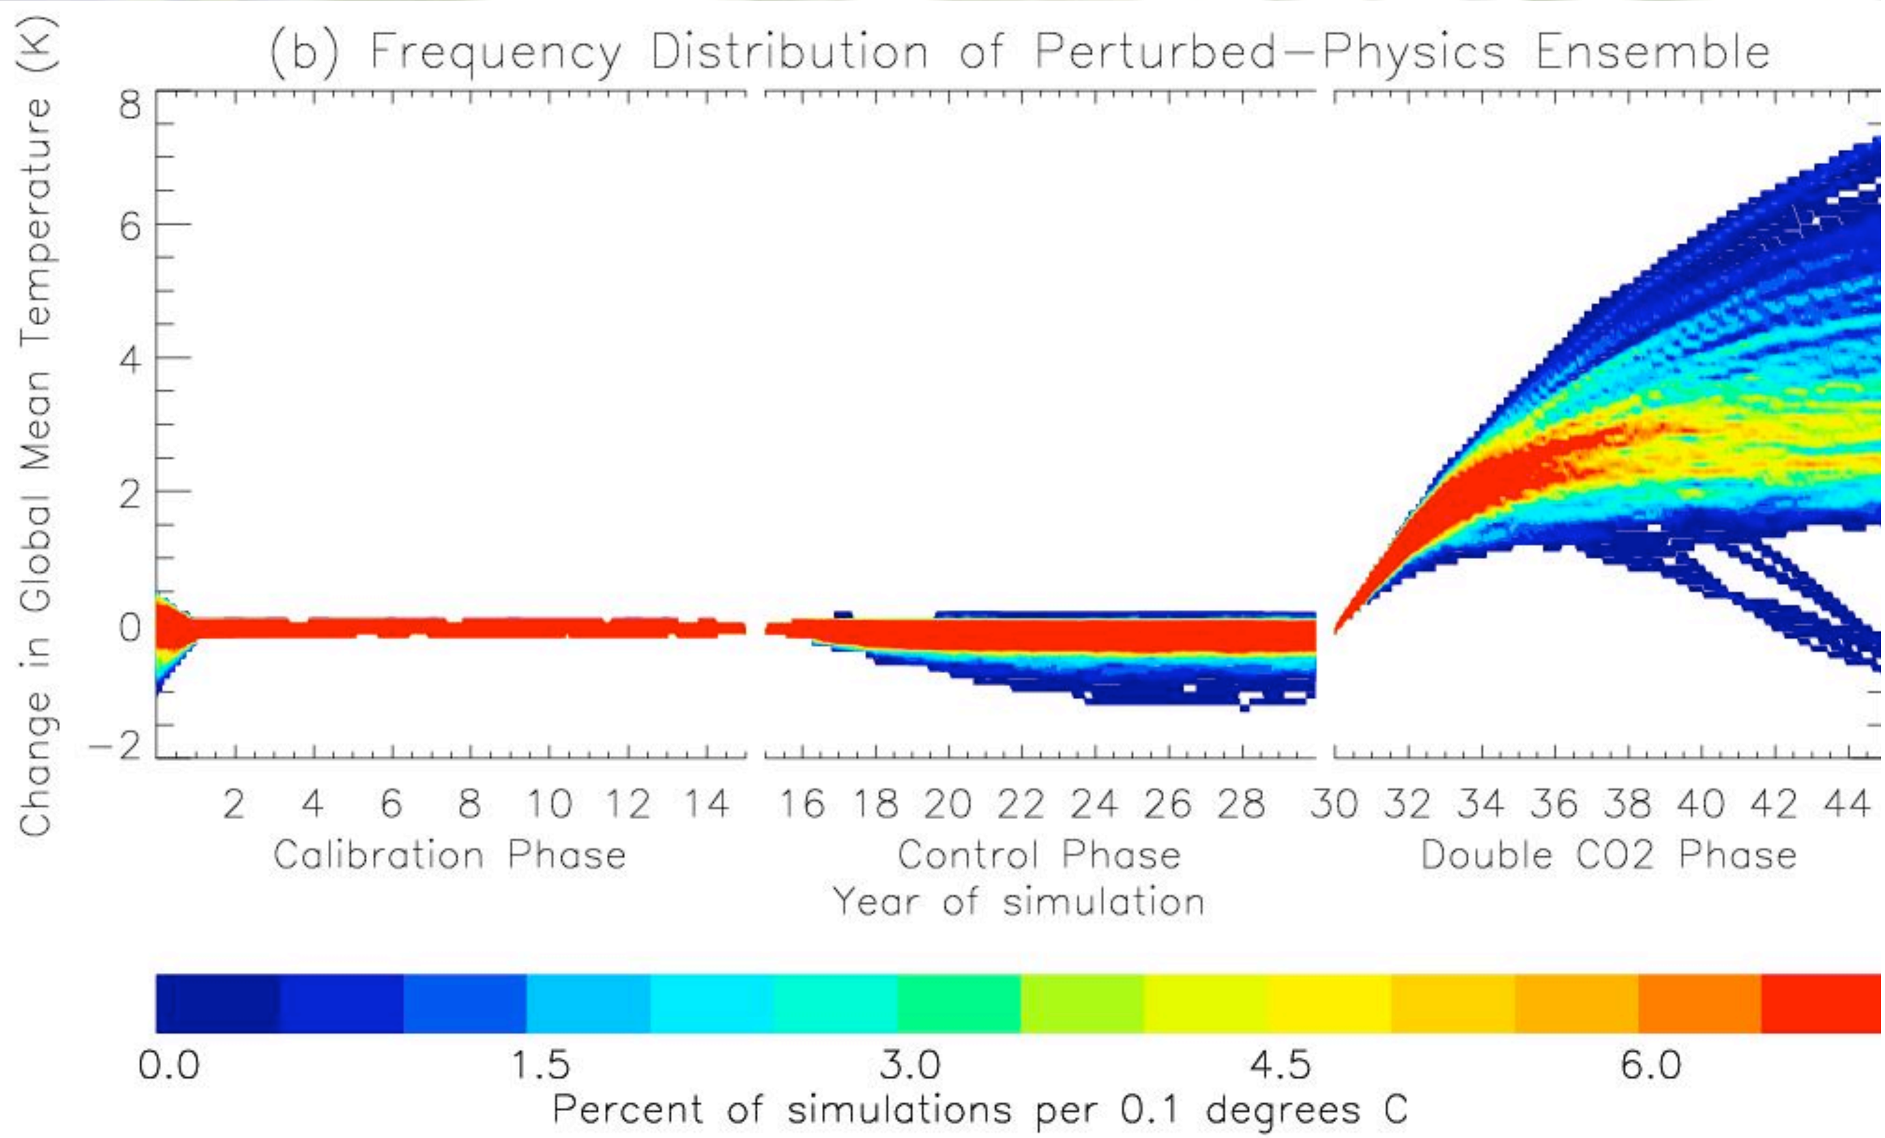

Un-physically strong low-cloud versus surface-heat-flux feedback in equatorial Pacific

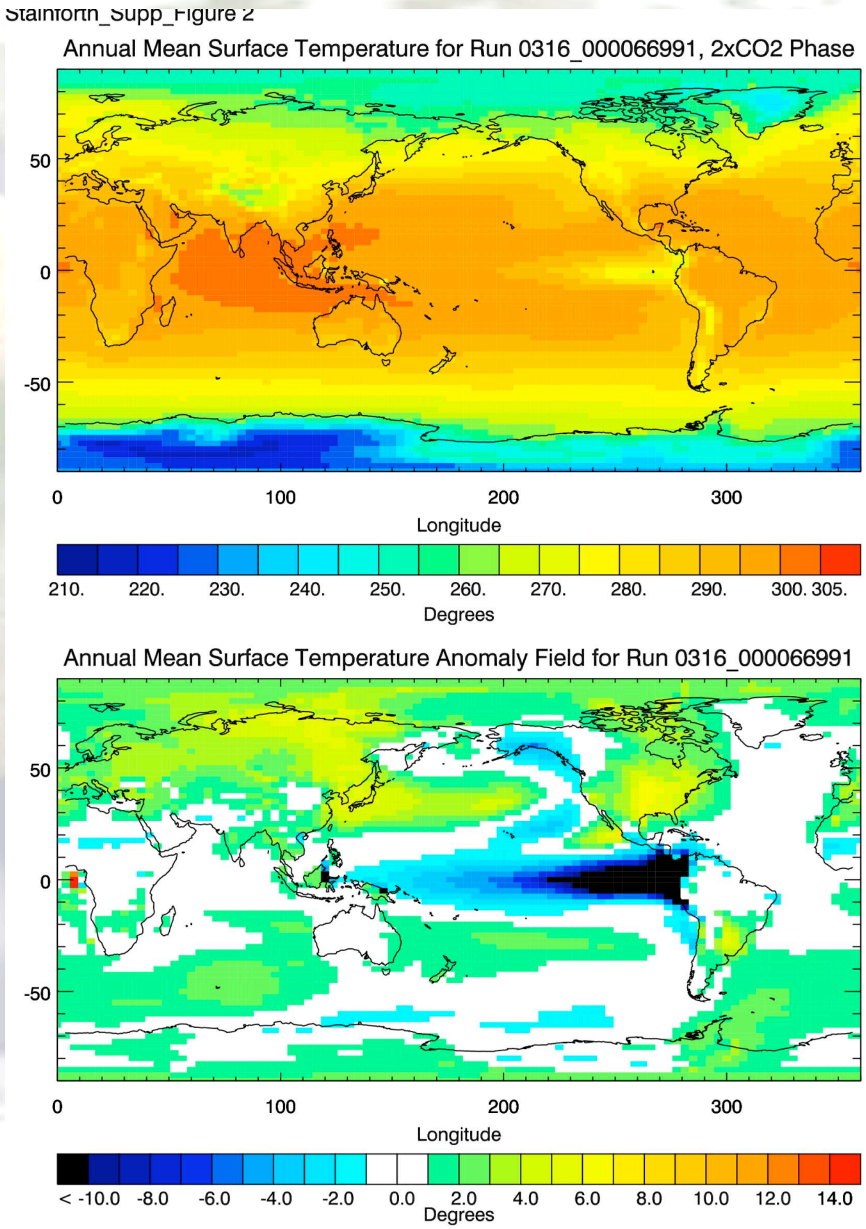

Do We Need So Many Simulations? Can we Predict Behaviour from a Smal Subset?

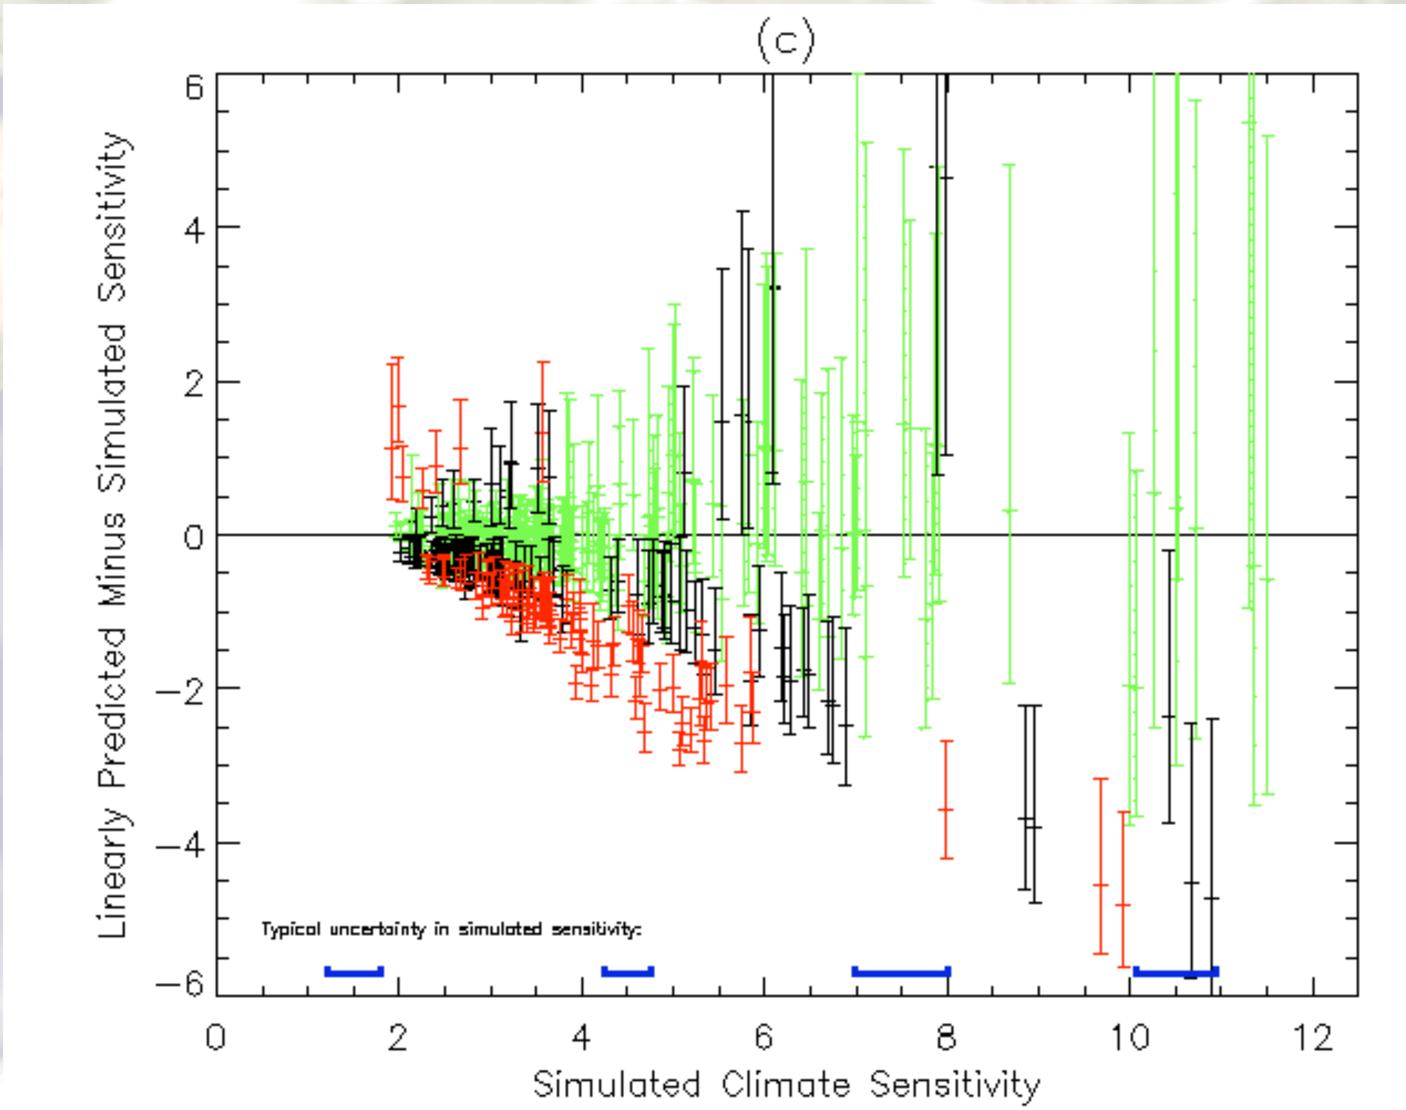

Climate sensitivities from climateprediction.net

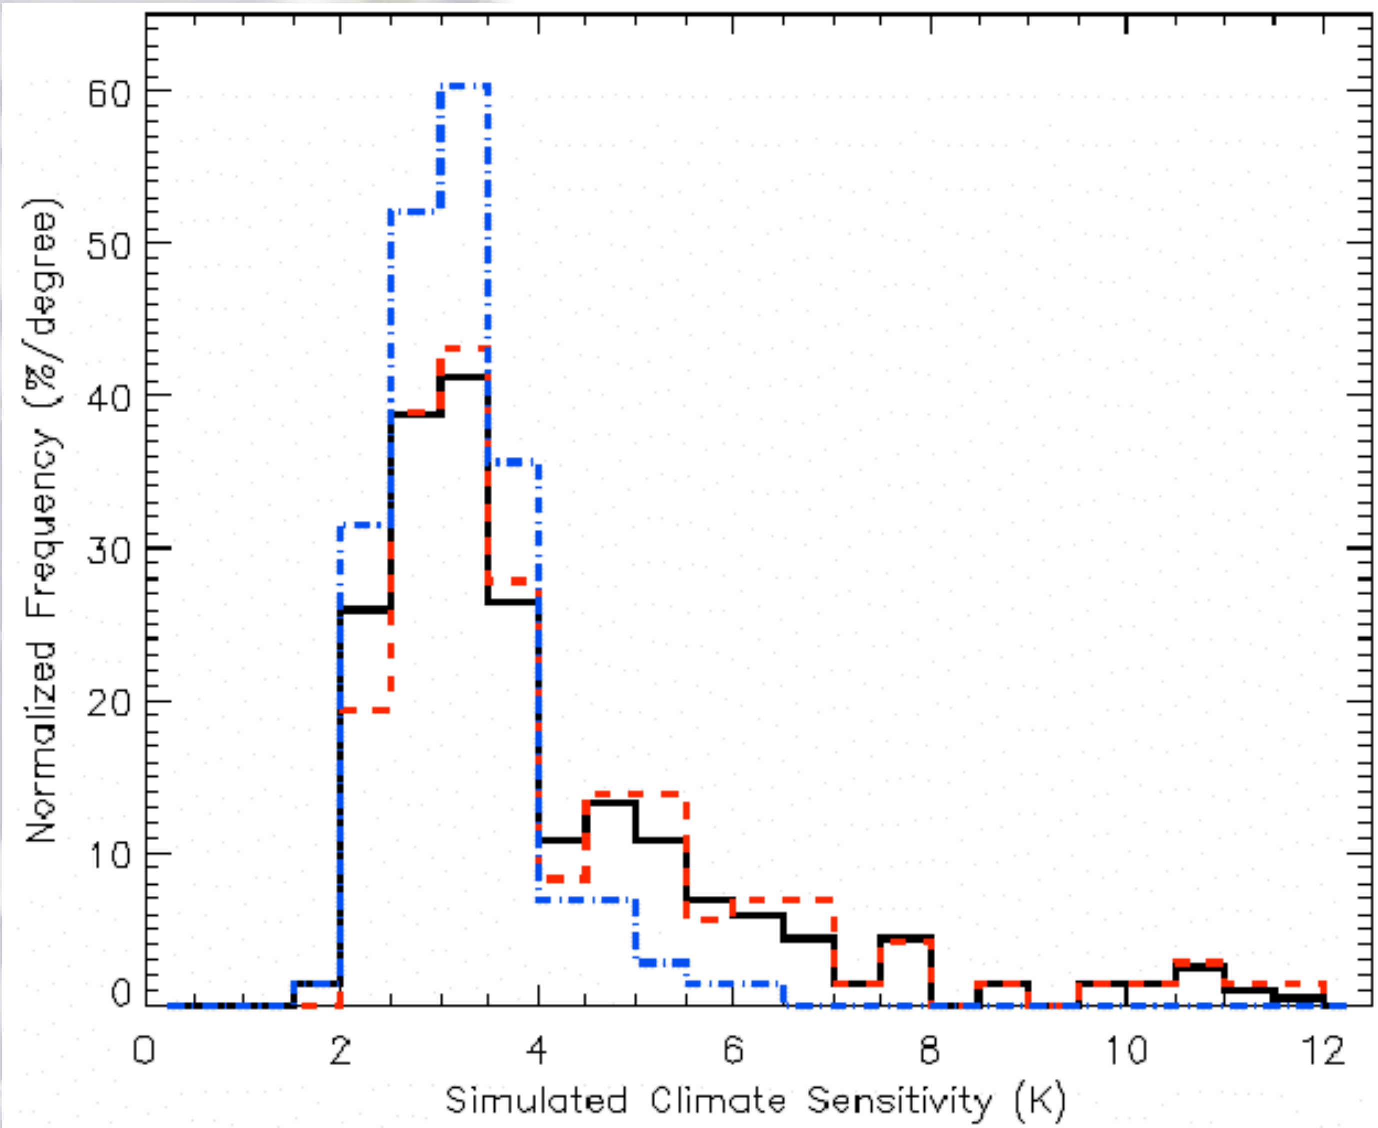

The frequency distribu
of simulated climate
sensitivity using all (2
model versions (black
model versions excep
those with perturbatio
the cloud-to-rain
conversion threshold
and all model version
except those with
perturbations to the
entrainment coefficier
(blue).

Sensitivity is the equil
response of the globa
mean temperature of
doubling atmospheric
levels of carbon dioxi

Stainforth et al, Nature, 27 Jan ‘05

And having got excited about the cold ones.

FREE

METRO

Check out our great going out guide

Thursday, January 27, 2005

www.metro.co.uk

Life, starts Page 17

Smoke billows into the sky above the port of Dudinka in Russia yesterday

11°C

That's how much hotter scientists believe the world will get ... and it will be worse in Britain

BY SUZY AUSTIN

THE world is likely to heat up by an average of 11°C by the end of the century, the biggest ever study of global warming showed yesterday.

And the effect could be even more marked in Britain, where temperatures could soar by up to 20°C unless greenhouse gases are cut.

Such a rise – far higher than the 2°C previously forecast – would see Britain endure tropical temperatures, flooding and devastating drought.

It would change the weather patterns of the world, melt the polar ice caps and warm the oceans, causing a surge in sea levels threatening the lives of billions of people.

The findings come from a study which tapped into the processing power of 100,000 home computers in 150 countries.

Researchers racked up the equivalent of 8,000 years of processing time as they ran 60,000 potential scenarios through the network, far more than the 128 scenarios the powerful computers at the Met Office can check in a year.

Each scenario was based on the assumption that carbon dioxide levels had reached double those of pre-Industrial Revolution times by the middle of this century.

Researcher David Stainforth, from Oxford University, said: 'An 11 degree warmed world would be a dramatically different world.

'Warming is not constant at all latitudes and tends to be greater at high latitudes.

'With a world warmed by 11 degrees there would be large areas of high latitude that could be 20 degrees warmer than they are today.

'I think it would probably not be a tropical paradise. The UK would be at the high end of this change, well into the teens as the temperature changes. I don't think we'll be building many snowmen in winter, or going sledging.'

The findings could mean world leaders need to toughen their commitment in the Kyoto agreements to cut CO_2 emissions to 5.2 per cent below 1990 levels by 2012.

The warning came as Tony Blair used the World Economic Forum in Switzerland to call for action on global warming and to pressure America to sign up to Kyoto.

Blair's call – Page 5

Stainforth et al, 2005, updated: raw distribution

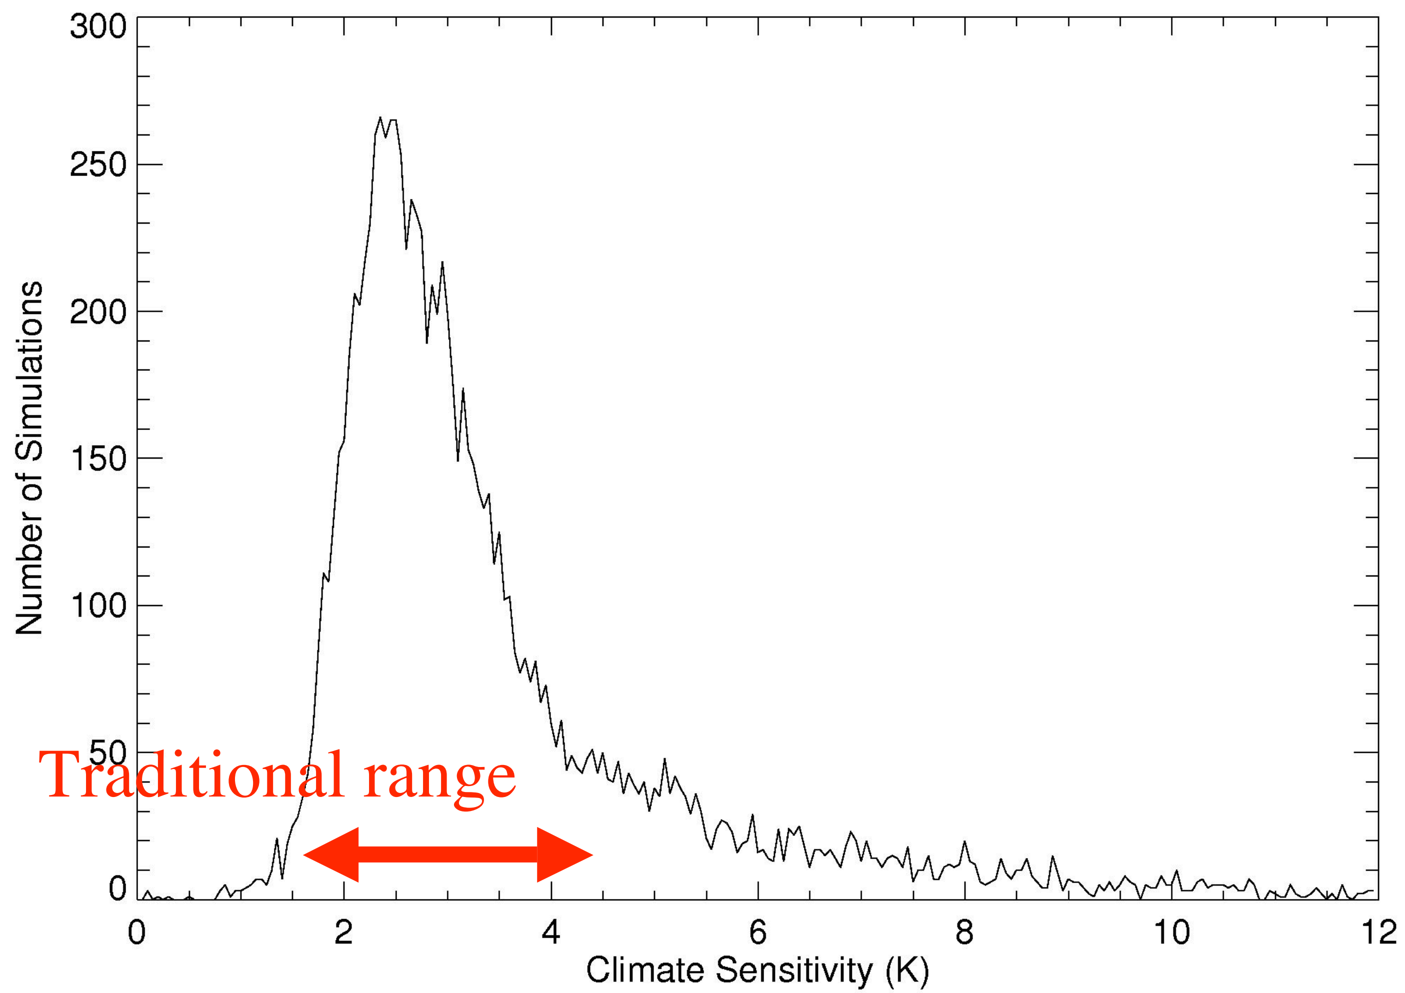

Raw histogram plotted against S^{-1}: a Gaussian with a lump on one side.

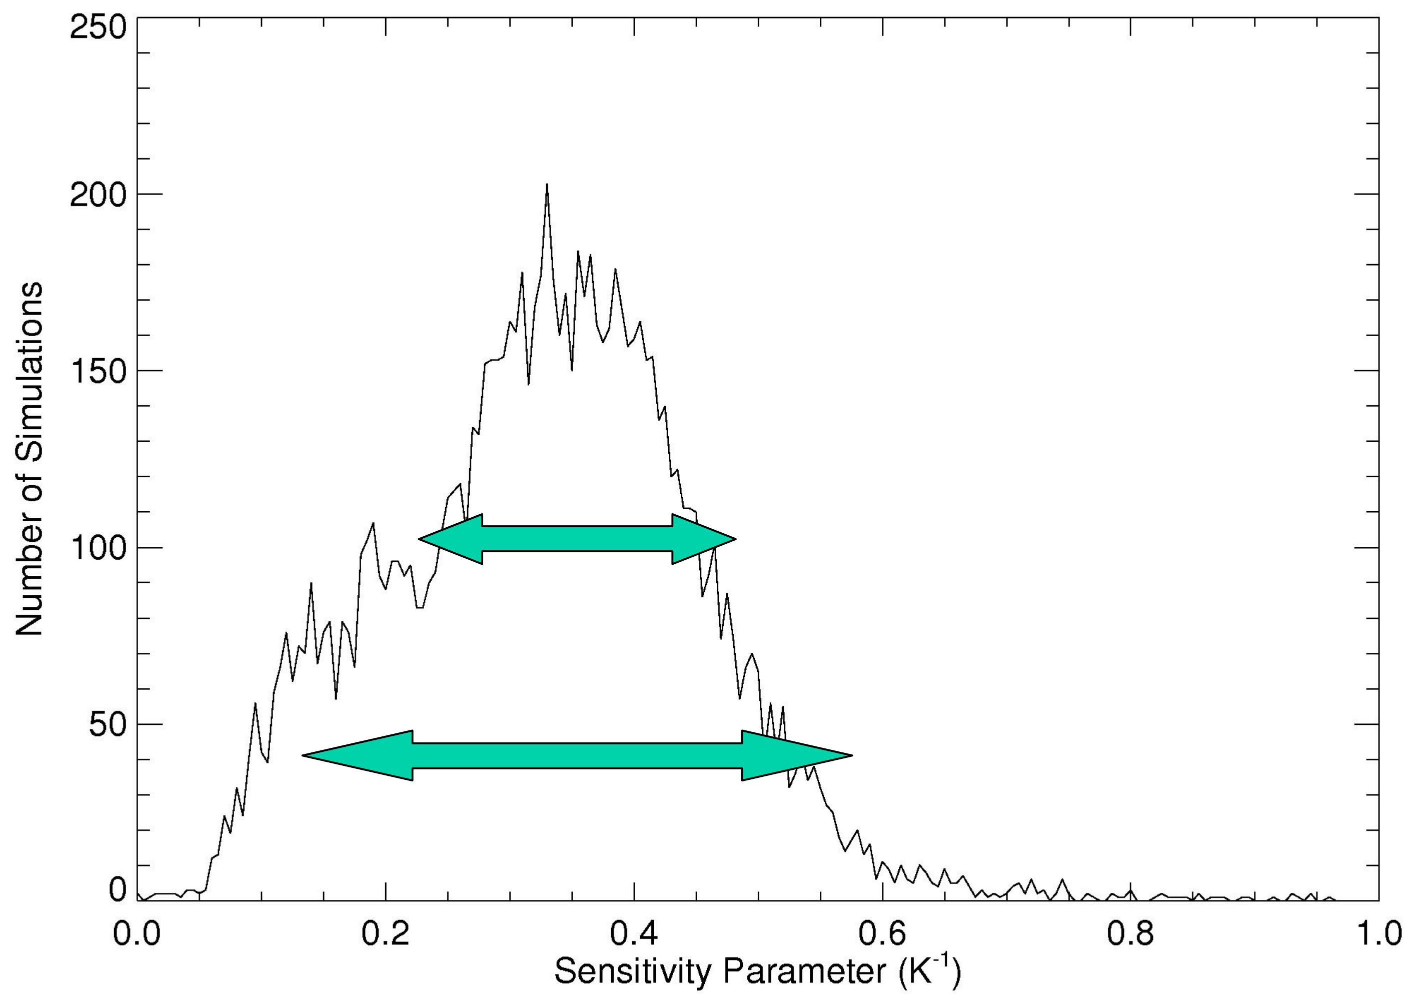

So why was anyone surprised by $S>10K$?

- Simple inspection of the distribution of feedback parameters, $\lambda=F_0/S$, suggests a Gaussian centred on

  $S^{-1}=0.35K^{-1}$ ($S=2.9K$) $\pm$ $0.2K^{-1}$ (2-σ) .
- This implies similar odds on obtaining a model with $S^{-1}>0.6K^{-1}$ ($S<1.7K$) as $S^{-1}<0.1K$ ($S>10K$).
- Would anyone have batted an eyelid if we had announced a small percentage of models with sensitivities $<1.7K$?

Many of these high sensitivity models will prove significantly less realistic than the original

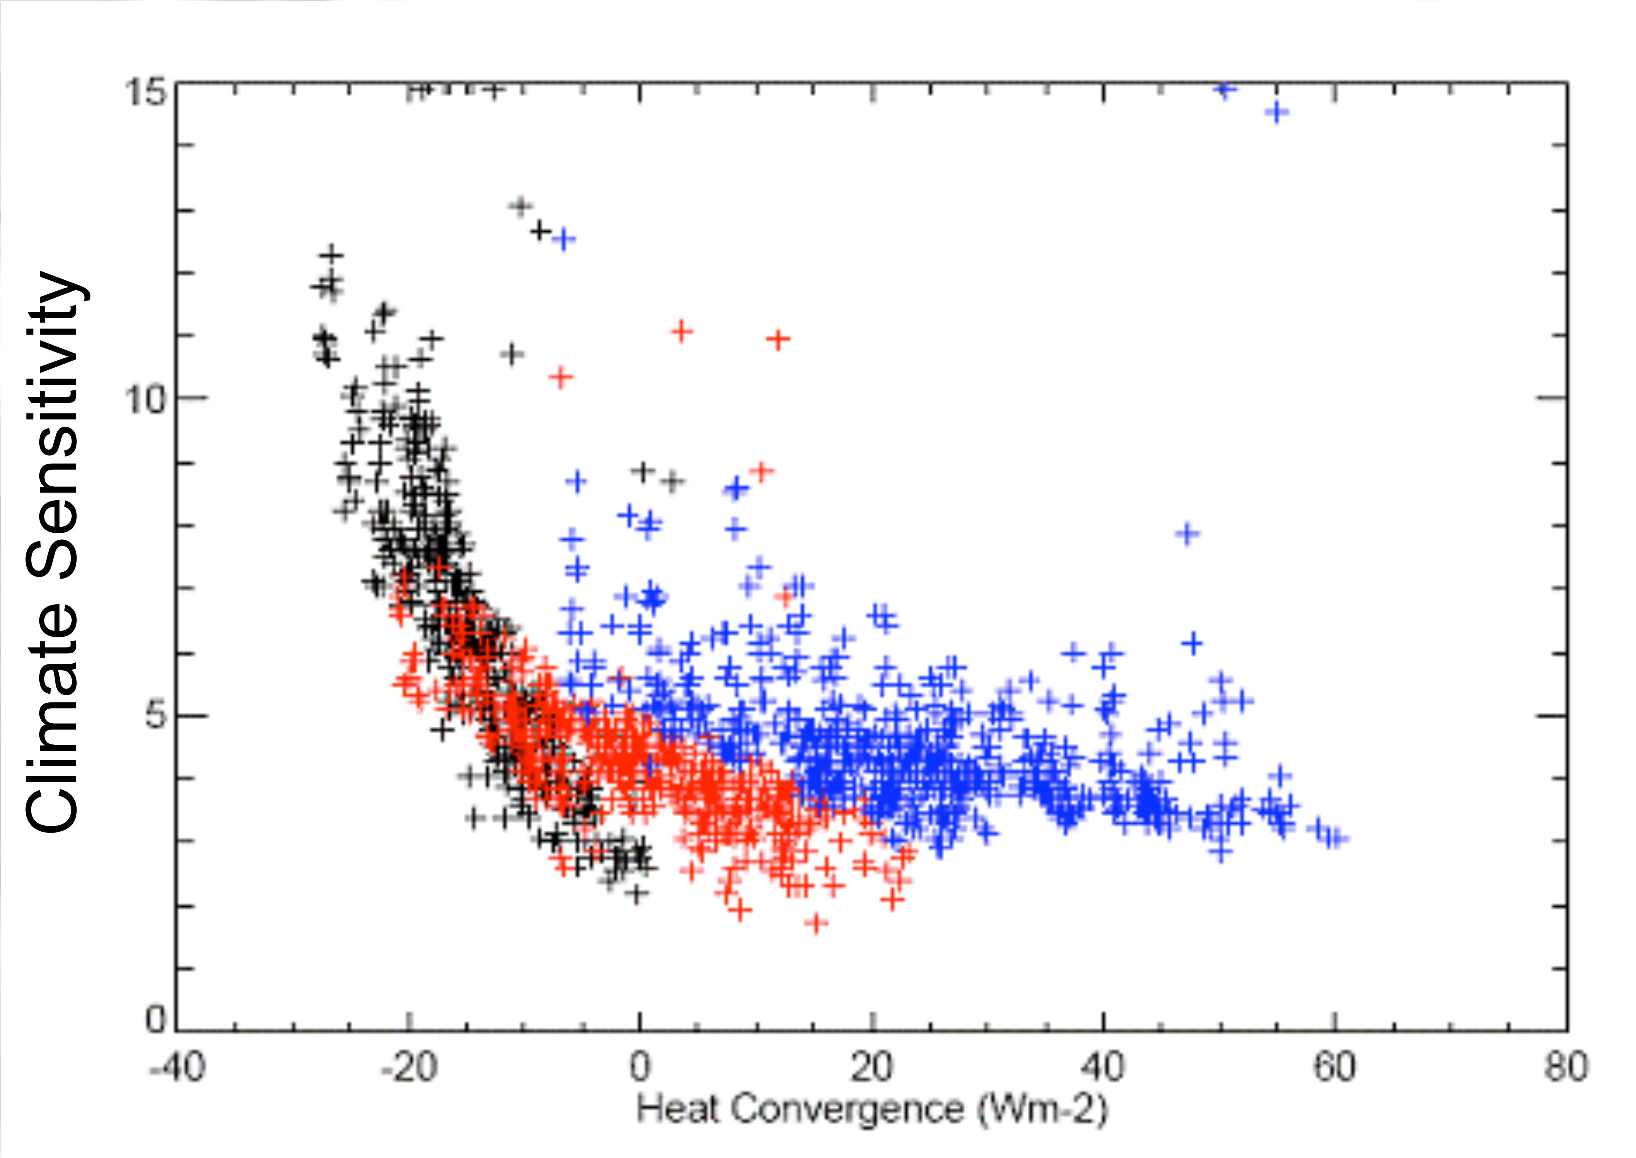

Global Top-of-Atmosphere Energy Imbalance

But not all

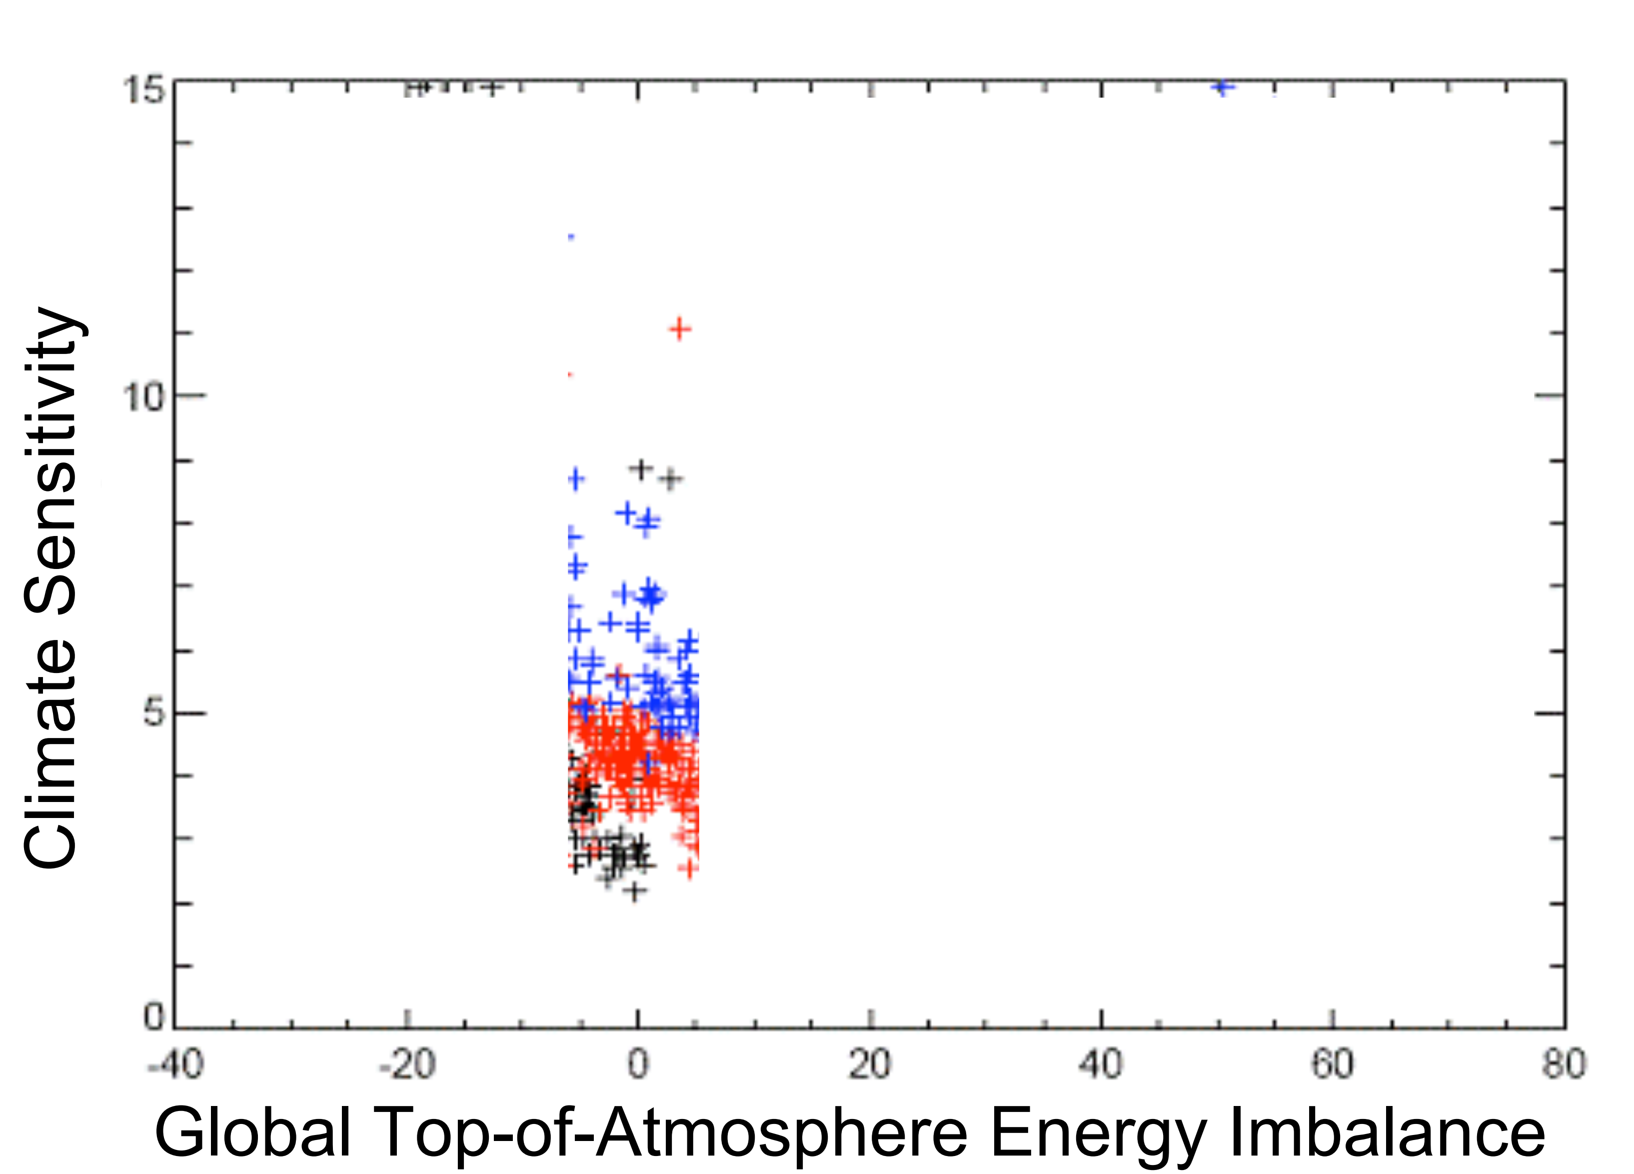
15
10
5
0
-40
-20
0
20
40
60
80
Climate Sensitivity
Global Top-of-Atmosphere Energy Imbalance

How can we estimate a distribution for S tha does not depend on arbitrary sampling?

- Histograms or weighted histograms depend to first order on sampling design.
- For example:
  - Choosing to weight different regions of parameter space by the “size” of parameter perturbations makes results depend on arbitrary definition of model parameters.
  - Choosing to sample S^{-1} uniformly over parameter intervals, or choosing to sample “forcing” uniformly i an energy budget analysis (Forster and Gregory), weights values of S by S^{-2} before any comparison with observations is made.
- Instead, we search for “Transfer Functions” between observable quantities and climate sensitivity

A transfer function approach to estimating climate sensitivity (Piani et al, 2005)

- We analyze a perturbed physics ensemble to identify a (vector) observable quantity that scales with the forecast quantity of interest (sensitivity).
- We then equate likelihood of observations given mode predictions of this observable with likelihood of the corresponding value for sensitivity.
- Although the models do not sample “model space” in any objective way, we hypothesize that the transfer function is invariant across models.
- Most importantly, this hypothesis can be easily tested by enlarging the ensemble or subdividing it.

Estimating λ from control climatology

Implied distribution allowing for scatter in transfer function

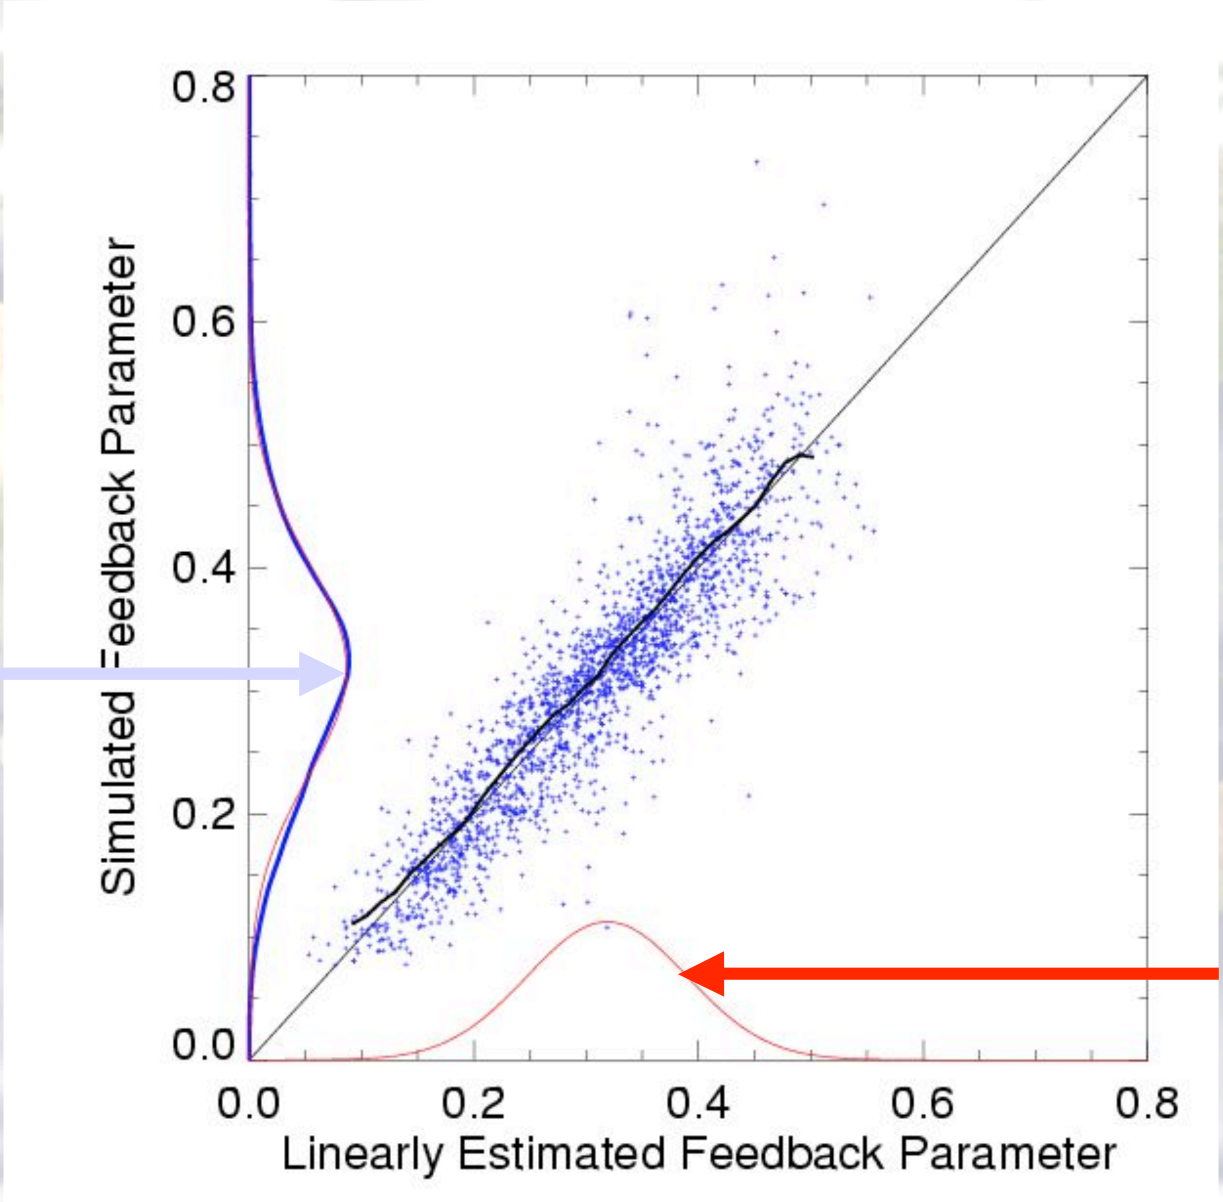

Best linear predictor o
λ/F_0 applie
to ERA-40

Estimating S from control climatology

Distribution from linear transfer function (wrong)

Distribution of S implied by actual transfer function

Likelihood function for λ/F_0 plotted against sensitivity

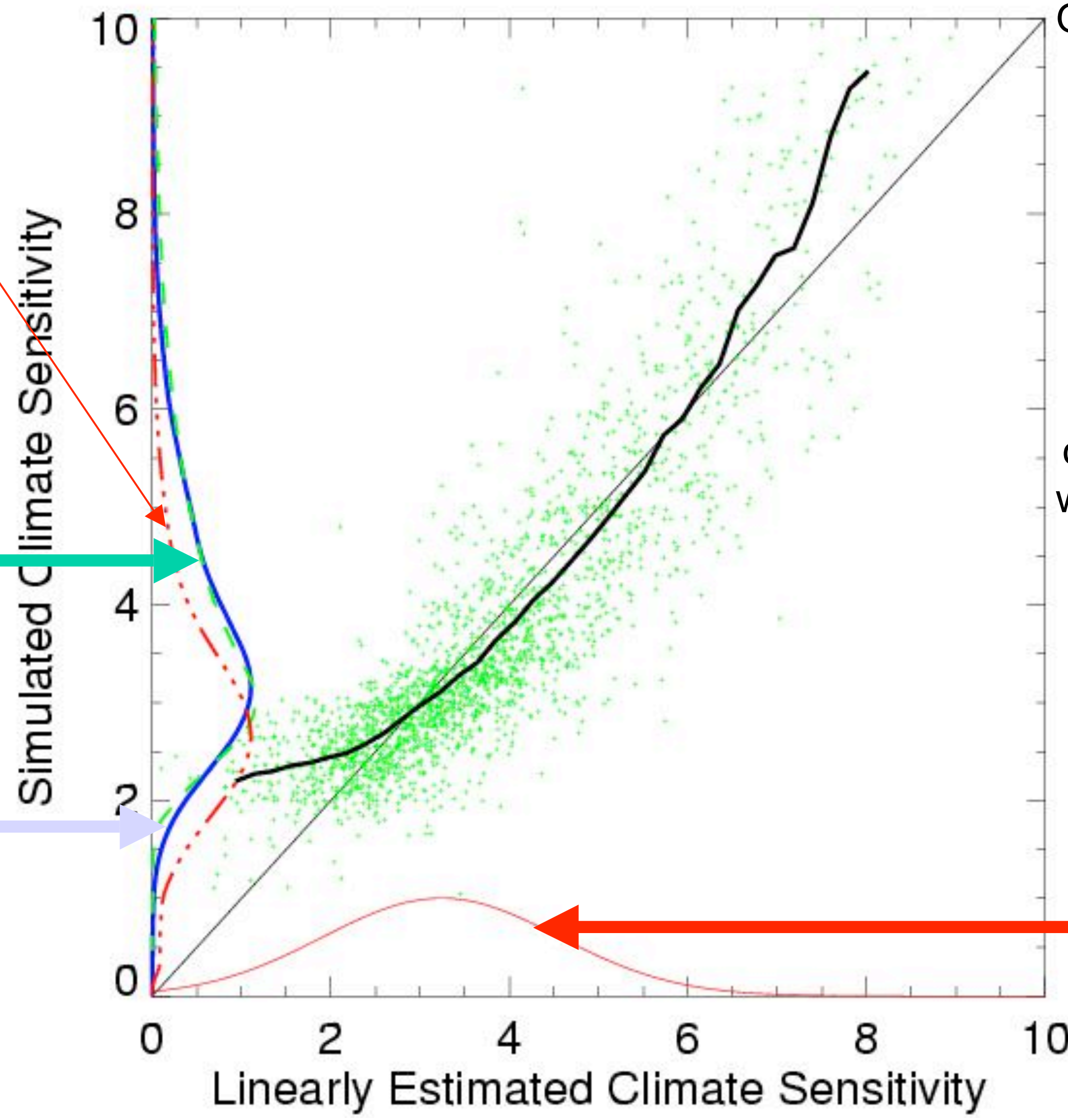

Objective: Develop a transfer functi identify a (vect observable quanti scales with the fo quantity of inte (sensitivity).

Our best estima climate sensitivity i with a 5-95% range 6.8K (we can't pin more extrem percentiles)

Best linea predictor applied to ERA-40

Piani et al, GRL, 12/2005

Implications of Piani et al

- Our best estimate of climate sensitivity is 3.3K, with a 5-95% range of 2.2-6.8K (we can’t pin down more extreme percentiles).
- Observable quantities scale with the feedback parameter and not with climate sensitivity - even if we search for predictors that are linear in sensitivity.
- So, if you want to estimate S via λ, you should estimat a likelihood function for λ and then plot it against S, without weighting by $d\lambda/dS=-F_0/S^2$.
- Results predicting S directly and via λ then agree.

Similar results using a neural network approach:

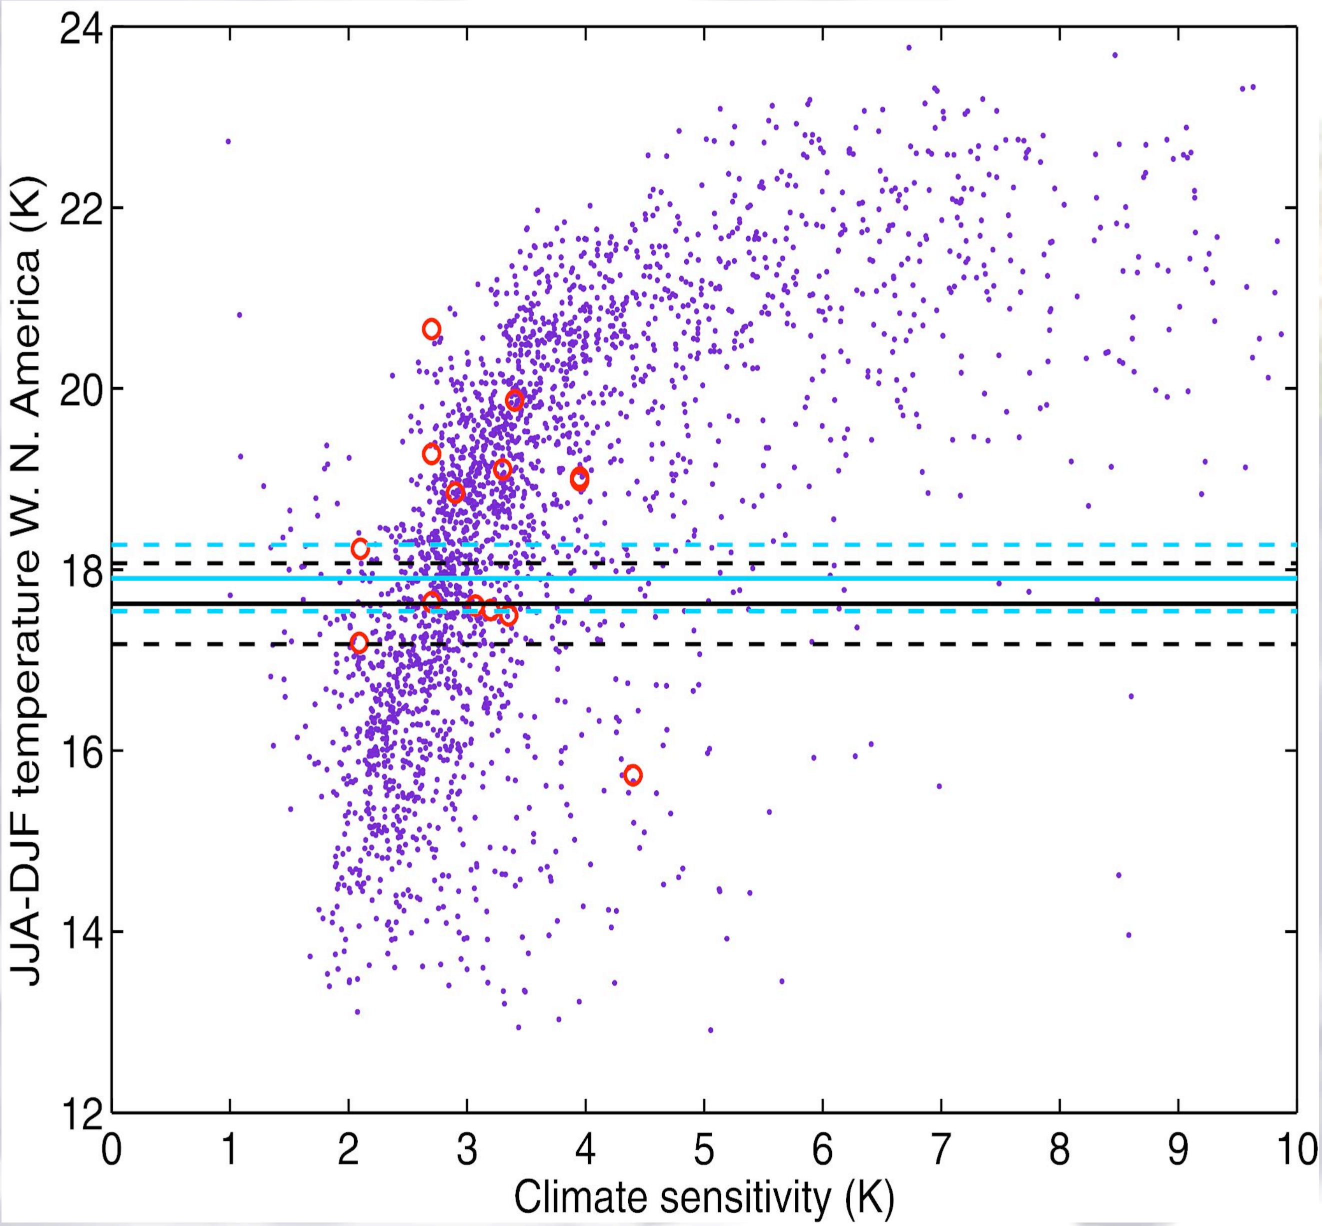

The amplitude of the seasonal c
in near surface temperature in
Western North America vs. clim
sensitivity. Each blue dot repres
a simulation from the
climateprediction.net project. Bl
horizontal lines mark the means
from the ERA40, NCEP and
HadCruT2v datasets over the sa
region.

The uncertainties in each of the
mean climatologies are smaller
the spread of the three datasets
Red circles mark amplitudes of
seasonal cycle in the IPCC 20th
century simulations.

NB: climateprediction.net
simulations represent orders of
magnitude increase in typical m
ensemble size.

Knutti et al, Journal of Climate, 2006

We were by no means the first to report the possibility of S>7K

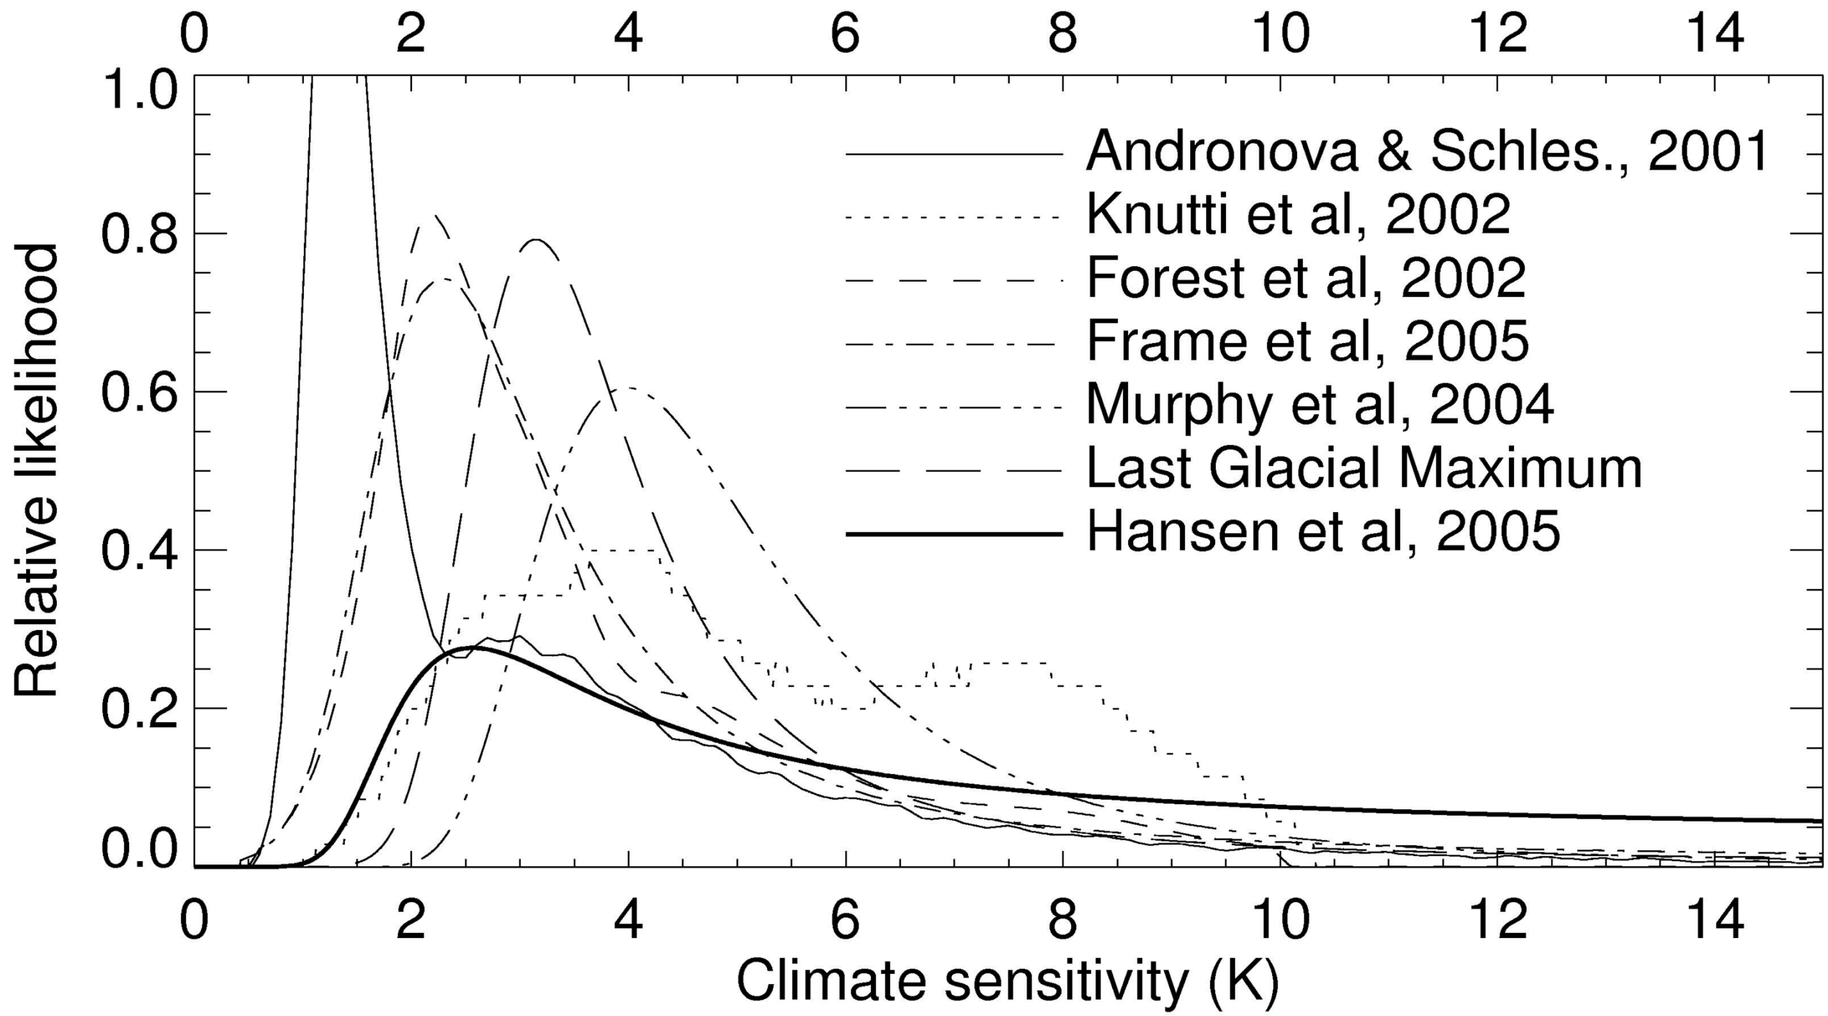

climate*prediction*.net

What’s next?

climate*prediction*.net

Next Step 1 – Regional Modelling

- NERC KT Project Using UKMO PRECIS model http://precis.metoffice.com/
- Currently 32-bit, Linux 100x100 regional model, max 20km grid cell resolution
- ‘embedded’ within our HadCM3L/BOINC project
- HadCM3L boundary conditions assimilated into PRECIS every model-day (unidirectional)
- Comparison of global versus regional resolution:

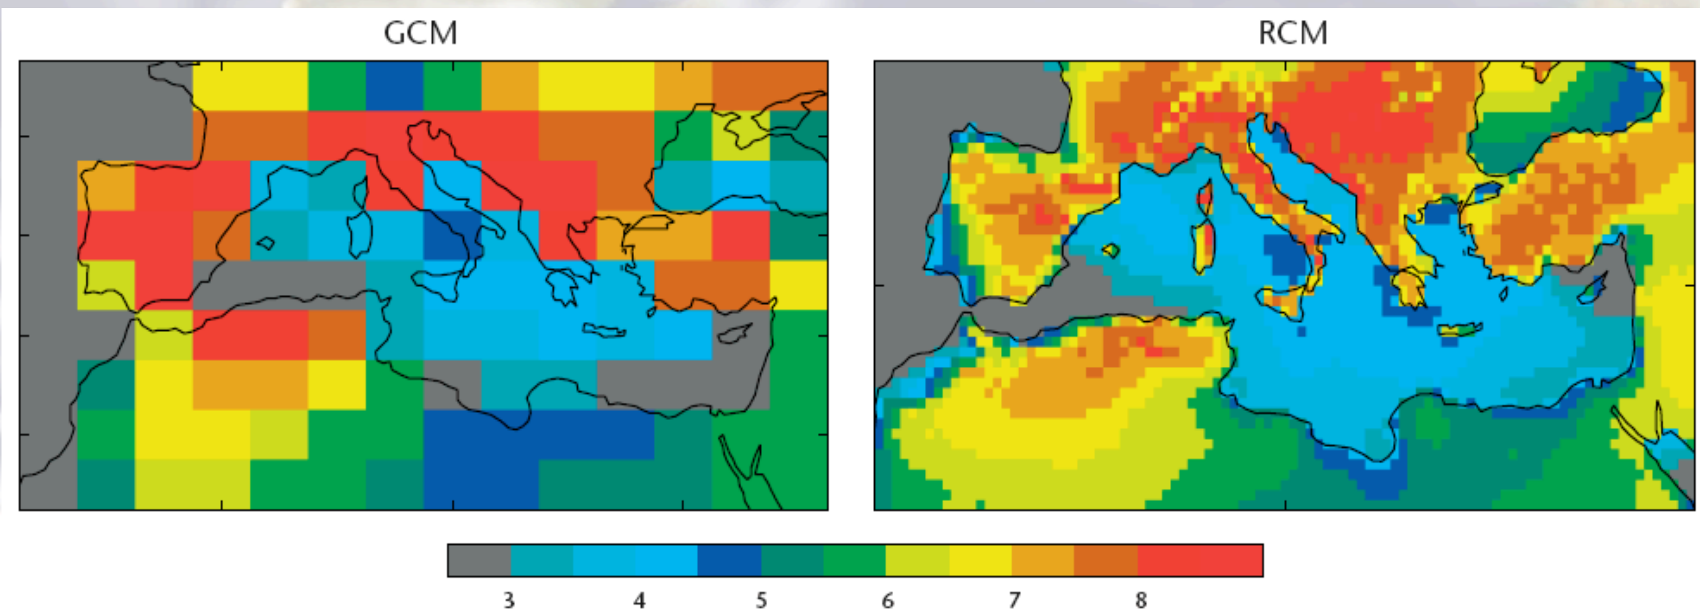

climate*prediction*.net

Next Step 2 – Higher Resolution Globa

- 100 km^2 at mid-latitudes – lg storm tracking, fronts etc
- Currently doing this at http://attribution.cpdn.org
- Recently have a proposal to do this with HadGAM1 (current generation MetOffice model)

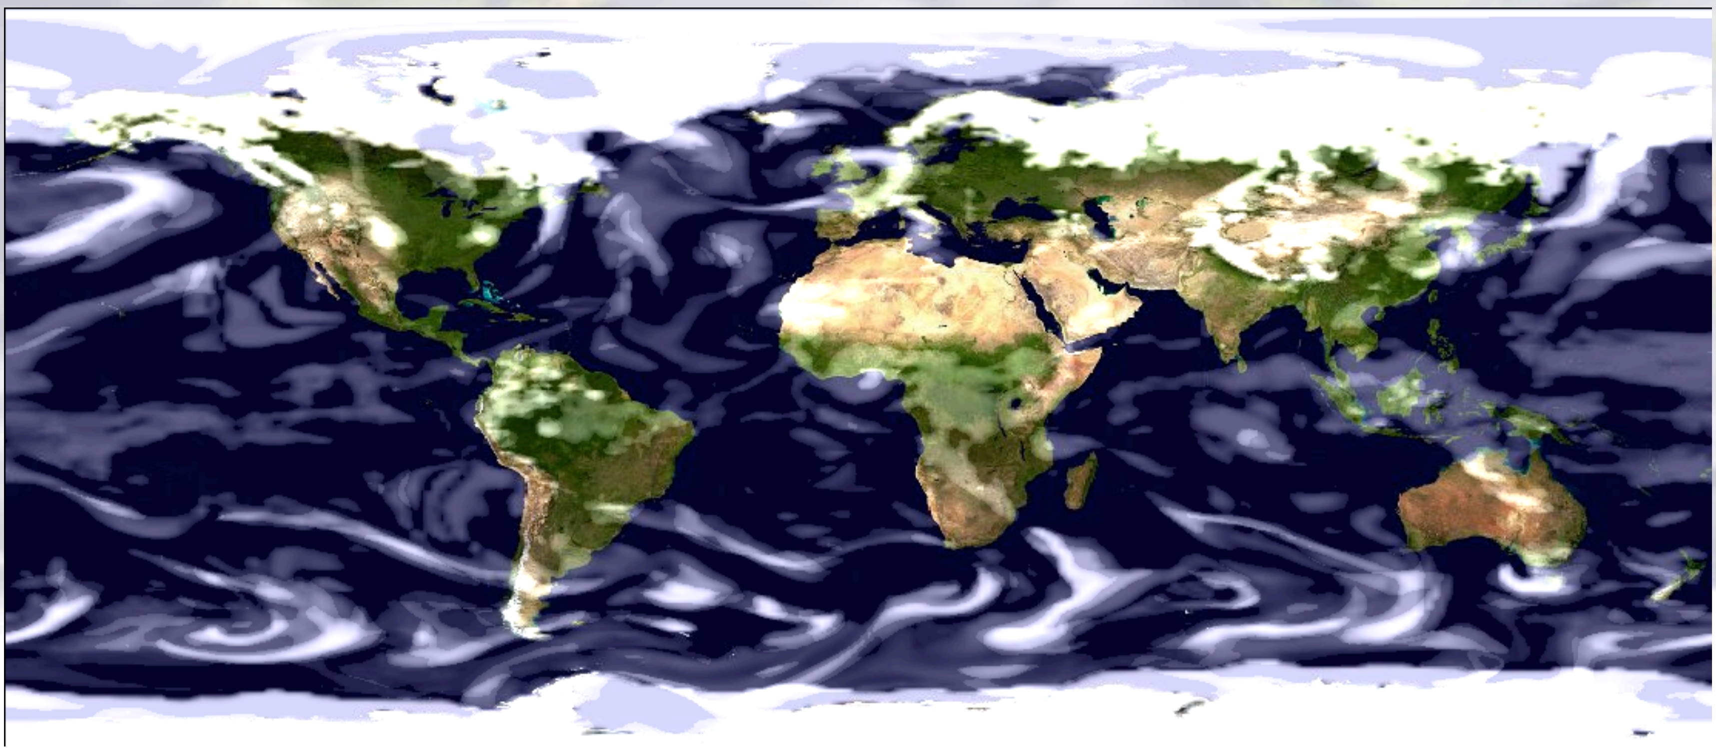

All of this would not be useful without observations!

Thank you.